

Patent Application for

**SOFTWARE-BASED NETWORK ATTACHED STORAGE SERVICES HOSTED ON
MASSIVELY DISTRIBUTED PARALLEL COMPUTING NETWORKS**

5

Inventors: Edward A. Hubbard

This application is a continuation-in-part application of the following co-pending applications: Application SN 09/538,543 entitled "DISTRIBUTED PARALLEL PROCESSING SYSTEM HAVING CAPABILITY-BASED INCENTIVES AND ASSOCIATED METHOD," Application SN 09/539,023
10 entitled "SWEEPSTAKES INCENTIVE MODEL AND ASSOCIATED SYSTEM," Application SN 09/539,448, entitled "CAPABILITY-BASED DISTRIBUTED PARALLEL PROCESING SYSTEM AND ASSOCIATED METHOD," Application SN 09/539,428 entitled "METHOD OF MANAGING DISTRIBUTED WORKLOADS AND ASSOCIATED SYSTEM," Application SN 09/539,106 entitled "NETWORK SITE TESTING METHOD AND ASSOCIATED SYSTEM," Application SN 09/538,542 entitled "NETWORK SITE CONTENT INDEXING METHOD AND ASSOCIATED SYSTEM," and Application SN 09/539,107 entitled "DISTRIBUTED BACK-UP SYSTEM AND ASSOCIATED METHOD," each of which was filed on March 30, 2000, and each of which is hereby incorporated by reference in its entirety.

This application is also a continuation-in-part application of the following co-pending applications: Application SN 09/603,740 entitled "METHOD OF MANAGING WORKLOADS AND ASSOCIATED DISTRIBUTED PROCESSING SYSTEM," Application SN 09/602,789 entitled "MACHINE GENERATED SWEEPSTAKES ENTRY MODEL AND ASSOCIATED DISTRIBUTED PROCESSING SYSTEM," Application SN 09/602,803 entitled "DATA SHARING AND FILE
25 DISTRIBUTION METHOD AND ASSOCIATED DISTRIBUTED PROCESSING SYSTEM," Application SN 09/602,844, entitled "DATA CONVERSION SERVICES AND ASSOCIATED DISTRIBUTED PROCESSING SYSTEM," and Application SN 09/602,983 entitled "CUSTOMER SERVICES AND ADVERTISING BASED UPON DEVICE ATTRIBUTES AND ASSOCIATED DISTRIBUTED PROCESSING SYSTEM," each of which was filed on June 23, 2000, and each of
30 which is hereby incorporated by reference in its entirety.

This application is also a continuation-in-part application of the following co-pending application: Application SN 09/648,832 entitled "SECURITY ARCHITECTURE FOR DISTRIBUTED PROCESSING SYSTEMS AND ASSOCIATED METHOD," which was filed on August 25, 2000, and which is hereby incorporated by reference in its entirety.

5

This application is also a continuation-in-part application of the following co-pending application: Application SN 09/794,969 entitled "SYSTEM AND METHOD FOR MONITIZING NETWORK CONNECTED USER BASES UTILIZING DISTRIBUTED PROCESSING SYSTEMS," which was filed on February 27, 2001, and which is hereby incorporated by reference in its entirety.

10

Technical Field of the Invention

This invention relates to distributed processing and more particularly to techniques and related methods for managing, facilitating and implementing distributed processing in a network environment.

Background

The NAS (Network Attached Storage) market is currently one of the fastest growing segments of the overall storage market. NAS devices typically take the form of stand-alone devices or systems that contain their own storage, processing, connectivity and management resources. Through these internal capabilities, the NAS device becomes a storage resource that connects to any network enabling many devices on the network to share data via the NAS device.

A NAS device is a term that is often used to refer to a specialized file server that connects to the network to provide storage capacity to network-connected users. NAS devices typically contain a slimmed-down (micro-kernel) operating system and file system. NAS devices often process only input/output (I/O) requests by supporting popular file sharing protocols such as NFS (UNIX) and SMB (DOS/Windows). NAS devices also typically use traditional LAN protocols such as Ethernet and TCP/IP to communicate over the network. Some general-purpose computers using full-blown operating systems such as Windows or UNIX are often labeled as NAS products because they can provide storage that is attached to a network. However, a pure NAS device is one built from scratch as a dedicated file I/O device. As used herein, the term "NAS device" broadly refers to a device that makes data storage resources available to network-connected user devices. A "dedicated NAS device" is a device whose primary operational purpose is for providing NAS services.

5 A related term is the storage area network (SAN). "SAN" is a term that is often used to refer to a back-end network connecting storage devices to servers via peripheral channels such as SCSI, SSA, ESCON and Fibre Channel. By using a SAN, the data storage traffic is off-loaded from the primary user network, which includes the user machines and the server machines. The servers then utilize the SAN to move and store data in storage devices connected to the storage network. SAN environments are becoming increasingly popular for enterprise storage solutions.

10 With respect to lower cost devices that are currently marketed as dedicated NAS devices, many have common personal computer (PC) specifications such as the following: Pentium III processor, 40GB HD and a 10/100 Ethernet connection, along with other components. Considering the specifications of an average corporate desktop computer, the capabilities of these desktop computers often exceed the capabilities computers being sold as dedicated NAS devices, and these desktop computers often sit idle for significant amounts of time each day.

Summary of the Invention

15 The present invention provides software-based network attached storage (NAS) services that are hosted on a massively distributed computing system. More particularly, a client agent program is configured to run on the network-connected client devices that are part of the distributed computing system to provide what appears to user devices as dedicated NAS functionality, for example, through the Internet. This NAS functionality takes advantage of unused or under-utilized resources of these distributed devices to advantageously provide NAS capabilities. In addition, storage priority controls may be utilized to facilitate the full use of available storage resources, such as user specified constraints, standard bit, block or file priority levels and direct bit, block or file priority markings.

25 In one embodiment, the present invention is a method of operating a distributed processing system to provide network attached storage (NAS) services, including configuring a client agent program to have a software-based NAS component and to take advantage of unused or under-utilized resources of a plurality of network-connected distributed devices put into operation for purposes distinct from operating the client agent program to provide NAS services, and utilizing the plurality of distributed devices running the client agent to provide NAS services to network-connected user devices with the network-connected user devices recognizing at least a portion of the plurality of distributed devices as dedicated NAS devices. In a more detailed respect, at least a portion of the plurality of network-

30

connected distributed devices may operate only as stand-alone dedicated NAS devices without knowledge of other NAS devices connected to the network. Still further, at least one of the network-connected distributed devices may store location information for data stored by the plurality of network-connected distributed devices. Still further the provided NAS services may be managed utilizing at least one centralized server system. And the server system may download the NAS component to the plurality of network-connected distributed devices with the NAS component being configured to operate as part of the client agent program already installed on the distributed devices. In yet another respect, the method of the present invention may include managing storage resources for the plurality of distributed devices using a storage priority control that facilitates full use of available storage resources. In addition, the storage priority control may include, for example, a parameter selectable through the client device, storage priority level schemes that prioritize data storage and deletion, and priority markings directly given to data or files.

In another embodiment, the present invention is a distributed processing system that provides network attached storage (NAS) services, including a client agent program and a plurality of network-connected distributed devices running the client agent program. The client agent program may be configured to have a software-based NAS component and to take advantage of unused or under-utilized resources of a network-connected distributed device put into operation for purposes distinct from operating the client agent program to provide NAS services. The plurality of network-connected distributed devices may operate to provide NAS services to network-connected user devices with the network-connected user devices recognizing at least a portion of the plurality of distributed devices as dedicated NAS devices. In more detailed respects, at least a portion of the plurality of network-connected distributed devices operate only as stand-alone dedicated NAS devices without knowledge of other NAS devices connected to the network. In addition, at least one of the network-connected distributed devices stores location information for data stored by the plurality of network-connected distributed devices. Still further, at least one centralized server system may be included to manage at least in part the NAS services provided by the plurality of network-connected distributed devices. And the NAS component may be configured to be downloaded from the server system to the plurality of network-connected distributed devices with the NAS component being configured to operate as part of the client agent program already installed on the distributed devices. In yet another embodiment, a storage priority control may be used to manage storage resources for the plurality of distributed devices with the storage priority control configured to facilitates full use of available storage resources.

Description of the Drawings

It is noted that the appended drawings illustrate only exemplary embodiments of the invention and are, therefore, not to be considered limiting of its scope, for the invention may admit to other equally effective embodiments.

5

FIG. 1A is a block diagram for a distributed processing system having client capability and incentive features, according to the present invention.

10

FIG. 1B is a block diagram for information flow among customer systems, server systems and client systems, according to the present invention.

FIG. 2A is a block diagram for a client system, according to the present invention.

FIG. 2B is a block diagram for processing elements within a client system, according to the present invention.

FIG. 2C is a block diagram for a client system agent installed on a client system, according to the present invention.

FIG. 2D is an example user interface for a client system agent, including incentive advertising, according to the present invention.

FIG. 3A is a block diagram for server systems, according to the present invention, including a control system, a sweepstakes system and a workload database.

25

FIG. 3B is a block diagram for server systems, customer systems, client systems and outsourced host systems, according to the present invention.

FIG. 3C is a block diagram for a server system processor, according to the present invention.

30

FIG. 3D is an alternative block diagram for a server system processor, according to the present invention.

FIG. 4 is a functional block diagram for an example sweepstakes incentive operation according to the present invention.

5 FIG. 5A is a block diagram for a distributed processing system for a network site indexing application, according to the present invention.

FIG. 5B is a functional block diagram for an indexing operation according to the present invention.

10 FIG. 6A is a block diagram for a server system according to the present invention, including a control system, a workload database, and a database of client capabilities balancing vectors.

FIG. 6B is a functional block diagram for client capabilities balancing of workloads according to the present invention.

FIG. 7A is a block diagram for a distributed processing system, according to the present invention, including example network sites on which site testing is to be conducted, such as load testing and/or quality-of-service (QoS) testing.

FIG. 7B is a functional block diagram for site-testing, according to the present invention.

FIG. 8 is a block diagram of a distributed processing system for a data backup application, according to the present invention.

25 FIG. 9 is a block diagram of an alternative representation of an interconnection fabric for a distributed processing system environment, according to the present invention.

30 FIG. 10 is a block diagram of a more detailed block diagram for a client system agent installed on a client system, according to the present invention.

FIG. 11A is a more detailed flow diagram for machine generated sweepstakes entries according to the present invention.

FIG. 11B is an alternative detailed flow diagram for machine generated sweepstakes entries according to the present invention.

5 FIG. 12A is a block diagram of a distributed processing system that allows customers to select client system attributes, according to the present invention.

FIG. 12B is a block flow diagram for client system attribute selection, according to the present invention.

10 FIG. 13A is a block diagram of a distributed processing system that provides data conversion services, according to the present invention.

FIG. 13B is a block flow diagram for data conversion services within a distributed processing system, according to the present invention.

FIG. 14A is a block diagram of a distributed processing system that provides data transmission caching, according to the present invention.

FIG. 14B is a block diagram of a distributed processing system that provides data sharing and file distribution, according to the present invention.

FIG. 15 is a block diagram of an alternative representation for a distributed processing system, according to the present invention.

25 FIG. 16 is a block diagram of a representation for a distributed processing system including security subsystems, according to the present invention.

30 FIG. 17A is a block diagram of a client system and server systems communication interface, according to the present invention.

FIG. 17B is a block diagram of communication layers for client system and server systems communication, according to the present invention.

FIG. 18A is a detailed block diagram for an embodiment of security activities for server systems, according to the present invention.

FIG. 18B is a detailed block diagram for an embodiment of security activities for client systems, according to the present invention.

FIG. 19 is a block diagram for a distributed processing system and environment in which network service providers are enabled to monetize their user bases.

FIG. 20 is a block diagram representing the components for a client agent along with a representative indication of responsibility for those components.

FIG. 21 is a block diagram for a network attached storage (NAS) embodiment for a distributed processing system according to the present invention.

Detailed Description of the Invention

The present invention contemplates the identification of the capabilities of distributed devices connected together through a wide variety of communication systems and networks and the aggregation of these capabilities to accomplish processing, storage, broadcasting or any other desired project objective. For example, distributed devices connected to each other through the Internet, an intranet network, a wireless network, home networks, or any other network may provide any of a number of useful capabilities to third parties once their respective capabilities are identified, organized, and managed for a desired task. These distributed devices may be connected personal computer systems (PCs), internet appliances, notebook computers, servers, storage devices, network attached storage (NAS) devices, wireless devices, hand-held devices, or any other computing device that has useful capabilities and is connected to a network in any manner. The present invention further contemplates providing an incentive, which may be based in part upon capabilities of the distributed devices, to encourage users and owners of the distributed devices to allow the capabilities of the distributed devices to be utilized in the distributed parallel processing system of the present invention.

The number of usable distributed devices contemplated by the present invention is preferably very large. Unlike a small local network environment, for example, as may be used by an Internet

Service Provider (ISP), which may include less than 100 interconnected computers systems to perform the tasks required by the ISP, the present invention preferably utilizes a multitude of widely distributed devices to provide a massively distributed processing system. With respect to the present invention, a multitude of distributed devices refers to greater than 1,000 different distributed devices. With respect to the present invention, widely distributed devices refers to a group of interconnected devices of which at least two are physically located at least 100 miles apart. With respect to the present invention, a massively distributed processing system is one that utilizes a multitude of widely distributed devices. The Internet is an example of a interconnected system that includes a multitude of widely distributed devices. An intranet system at a large corporation is an example of an interconnected system that includes a multitude of distributed devices, and if multiple corporate sites are involved, may include a multitude of widely distributed devices. A distributed processing system according to the present invention that utilizes such a multitude of widely distributed devices, as are available on the Internet or in a large corporate intranet, is a massively distributed processing system according to the present invention.

FIG. 1A is a block diagram for a distributed parallel processing system 100 according to the present invention. The network 102 is shown having a cloud outline to indicate the unlimited and widely varying nature of the network and of attached client types. For example, the network 102 may be the Internet, an internal company intranet, a local area network (LAN), a wide area network (WAN), a wireless network, a home network or any other system that connects together multiple systems and devices. In addition, network 102 may include any of these types of connectivity systems by themselves or in combination, for example, computer systems on a company intranet connected to computer systems on the Internet.

FIG. 1A also shows client systems 108, 110 ... 112 connected to the network 102 through communication links 118, 120 ... 122, respectively. In addition, server systems 104, other systems 106, and customer systems 152 are connected to the network 102 through communication links 114, 116 and 119, respectively. The client system capabilities block 124 is a subset of the server systems 104 and represents a determination of the capabilities of the client systems 108, 110 ... 112. The incentives block 126 is also a subset of the server systems 104 and represents an incentive provided to the users or owners of the clients systems 108, 110 ... 112 for allowing capabilities of the clients systems 108, 110 ... 112 to be utilized by the distributed processing system 100. The client systems 108, 110 and 112 represent any number of systems and/or devices that may be identified, organized and utilized by the server systems

104 to accomplish a desired task, for example, personal computer systems (PCs), internet appliances, notebook computers, servers, storage devices, network attached storage (NAS) devices, wireless devices, hand-held devices, or any other computing device that has useful capabilities and is connected to a network in any manner. The server systems 104 represent any number of processing systems that provide the function of identifying, organizing and utilizing the client systems to achieve the desired tasks.

The incentives provided by the incentives block 126 may be any desired incentive. For example, the incentive may be a sweepstakes in which entries are given to client systems 108, 110 ... 112 that are signed up to be utilized by the distributed processing system 100. Other example incentives are reward systems, such as airline frequent-flyer miles, purchase credits and vouchers, payments of money, monetary prizes, property prizes, free trips, time-share rentals, cruises, connectivity services, free or reduced cost Internet access, domain name hosting, mail accounts, participation in significant research projects, achievement of personal goals, or any other desired incentive or reward.

As indicated above, any number of other systems may also be connected to the network 102. The element 106, therefore, represents any number of a variety of other systems that may be connected to the network 102. The other systems 106 may include ISPs, web servers, university computer systems, and any other distributed device connected to the network 102, for example, personal computer systems (PCs), internet appliances, notebook computers, servers, storage devices, network attached storage (NAS) devices, wireless devices, hand-held devices, or any other connected computing device that has useful capabilities and is connected to a network in any manner. The customer systems 152 represents customers that have projects for the distributed processing system, as further described with respect to FIG. 1B. The customer systems 152 connect to the network 102 through the communication link 119.

It is noted that the communication links 114, 116, 118, 119, 120 and 122 may allow for communication to occur, if desired, between any of the systems connected to the network 102. For example, client systems 108, 110 ... 112 may communicate directly with each other in peer-to-peer type communications. It is further noted that the communication links 114, 116, 118, 119, 120 and 122 may be any desired technique for connecting into any portion of the network 102, such as, Ethernet connections, wireless connections, ISDN connections, DSL connections, modem dial-up connections, cable modem connections, fiber optic connections, direct T1 or T3 connections, routers, portal computers, as well as any other network or communication connection. It is also noted that there are any

number of possible configurations for the connections for network 102, according to the present invention. The client system 108 may be, for example, an individual personal computer located in someone's home and may be connected to the Internet through an Internet Service Provider (ISP). Client system 108 may also be a personal computer located on an employee's desk at a company that is connected to an intranet through a network router and then connected to the Internet through a second router or portal computer. Client system 108 may further be personal computers connected to a company's intranet, and the server systems 104 may also be connected to that same intranet. In short, a wide variety of network environments are contemplated by the present invention on which a large number of potential client systems are connected.

FIG. 1B is a block diagram for information flow 150 among customer systems 152, server systems 104 and client system 134, according to the present invention. The server systems 104, as discussed above, may include any number of different subsystems or components, as desired, including client system capabilities block 124 and incentives block 126. The server systems 104 send project and benchmark workloads 130 to client systems 134. A benchmark workload refers to a standard workload that may be used to determine the relative capabilities of the client systems 134. A project workload refers to a workload for a given project that is desired to be completed. The project workload may be, for example, a workload for projects such as network site content indexing, network site testing including network site load testing and network site quality of service testing, data back-up, drug design, drug interaction research, chemical reaction studies, bioinformatics including genetic and biological analyses, human genome analyses, pair-wise comparisons including fingerprint and DNA analyses, data mining, internet hosting services, intranet hosting services, auction services, market clearing services, payment systems, bioinformatic simulations, knowledge management services, trading services, data matching services, graphics rendering, or any other desired project.

Client systems 134, as discussed above, may be any number of different systems that are connected to the server systems 104 through a network 102, such as client systems 108, 110 ... 112 in FIG. 1A. The client systems 134 send results 132 back to the server systems 104 after the client systems 134 complete processing any given workload. Depending upon the workload project, the server systems 104 may then provide results 156 to customer systems 152. The customer systems 152 may be, for example, an entity that desires a given project to be undertaken, and if so, provides the project details and data 158 to the server systems 104.

FIG. 2A is a block diagram for an example client system 108 according to the present invention. In this simplified block diagram, an original workload 204 is received through line 208 from an interface 206. The original workload 204 represents a portion of the processing, storage or other activity required to complete the desired task for which the server system 104 is trying to accomplish. This original workload 204 is sent by the server system 104 through the network 102 and received by the client system 108 through communication link 118. The client system 108 processes the original workload 204. Following line 212, results 202 are then stored for transferring along line 210 to interface 206. Interface 206 may then communicate the results back to the server system 104 through communication line 118, or to other client systems (for example, with peering of client systems) and then through the network 102.

It is noted that the workload received by client system 108 and the processing or activity performed may depend up a variety of factors, as discussed further below. In part, this workload allocated by the server system 104 to each client system 108, 110 and 112 may depend upon the capabilities of the client system, such as the processing power, disk storage capacity, communications types, and other capabilities available from the various components of the systems within the client system 108.

The server systems 104 can select the workloads for the client system 108 and may control when these workloads are performed, through operational code (i.e., an agent) residing and installed on the client system 108. Alternatively, the owner or user of the client system 108 may determine when workloads are procured or obtained from the server systems 104, as well as when these workloads are performed, for example, by accessing the server systems 104 through the network 102. For example, the server systems 104 may download to the client system 108 upon request one or more workloads. At the same time, an agent residing on the client system 108 may operate to process the workload or multiple workloads downloaded to the client system 108. It is noted, therefore, that the agent may be simultaneously managing more than one workload for any number of projects. When the workload is complete, the agent may inform the owner or user of the client system 108 the results are ready to be communicated back. The client system 108 may then upload results to the server system 104 and download new workloads, if desired. Alternatively, these logistical and operational interactions may take place automatically through control of the agent and/or the server systems 104.

FIG. 2B is a block diagram for processing elements within a client system 108 according to the present invention. In this diagram, client system 108 is contemplated as a personal computer. In a

personal computer, an internal bus 260 would typically have a variety of different devices connected to it. For example, a CPU 250 could be connected through the bus 260 to a video processor 252, a floating point processor 254 (often integrated within the CPU itself), and digital signal processors 256 (DSPs), such as those found on sound cards and modems. In addition, any of a variety of other processing devices 258 may be included. Furthermore, other types of devices may be connected, such as hard drives 264, which provide disk storage capabilities, and a digital camera 262.

It is noted, therefore, that the capabilities for client systems 108, 110 ... 112 may span the entire range of possible computing, processing, storage and other subsystems or devices that are connected to a system connected to the network 102. For example, these subsystems or devices may include: central processing units (CPUs), digital signal processors (DSPs), graphics processing engines (GPEs), hard drives (HDs), memory (MEM), audio subsystems (ASs), communications subsystems (CSs), removable media types (RMs), and other accessories with potentially useful unused capabilities (OAs). In short, for any given computer system connected to a network 102, there exists a variety of capabilities that may be utilized by that system to accomplish its direct tasks. At any given time, however, only a fraction of these capabilities are typically used on the client systems 108, 110 ... 112. The present invention can take advantage of these unused capabilities.

It is also noted that along with receiving the workload, the client system 108 will also receive an agent that manages the completion of the workload. This agent may be software that is customized for the particular computer system and processing capabilities of the client system 108. For example, if the client system is a personal computer as shown in FIG. 2B, the agent may be a program that operates in the background of the computer's operating system. When the agent determines that there is unused processing or other capabilities, the agent may take advantage of it. For example, if the user is using a word processing application to create a document, little processing power is being utilized by the word processing program, leaving the computer's CPU and video processor underutilized. Thus, the agent could execute commands to these processors during dead cycles. In this way, the agent may facilitate the completion of workload processing in a reduced time. In addition, this agent may be self-updating upon connecting to the server systems 104, so that the agent may be kept up to date with current software revisions and workload activities. It is also noted that the agent may manage work on multiple workloads at the same time, so that any given distributed device connected to the network 102 may be working on a plurality of workloads at any given time.

FIG. 2C is a block diagram for an example client system agent 270. The agent 270 may include a security subsystem 272 that controls the interface of the client system 108 with the agent 270. The security subsystem 272 may help keep the workloads secure and may help to keep the client systems 108 from suffering any security problems in completing the workload. For example, the agent 272 may operate to keep viruses from attacking the client system 108 while the client system 108 is processing the workload through the operation of the agent. The security subsystem 272, therefore, may provide the interface for the workloads 130 and the results 132.

The clients system agent 270 may also include a workload engine 274, a statistics / user interface / incentive advertising block 276, and a workload package and update processing block 278. In the example shown in FIG. 2C, workloads 130 pass through the security subsystem 272 and along line 280 to the workload package and update processing block 278. In this block 278, the agent 270 may be updated by the server systems 104. Alternatively, the agent 270 may determine, when connected to the server systems 104, whether it needs to be updated and then accomplish that updating automatically. Once the workload package is processed, the workload engine 274 may receive the workload following line 288. The workload engine 274 works on the workload, ultimately completing the workload. The results or status of the workload may then be sent through the security subsystem 272 following line 282. The results 132 may then be provided back to the server systems 104.

The statistics / user interface / incentive advertising block 276 may provide workload, incentive and other statistics, as well as any other desired interface features, to the user of the client system. For example, the block 276 may show a user the expected amount of processing time it will take for the client system to complete a workload task based upon the capabilities of the system. As also shown, the block 276 may receive information following lines 286 and 284 from the workload package and update processing block 278 and from the workload engine 274. If desired, security information from the security subsystem 272 could also be displayed to the user of the client system. It is noted that the information displayed to the user of the client system may be modified and selected as desired without departing from the present invention.

With respect to incentive advertising, the block 276 may also show the user of the client system how this processing time might change depending upon various possible upgrades to the capabilities of the client system, such as a faster microprocessor, more memory, more disk storage space, etc. Furthermore, the client system capabilities may be shown correlated to the incentives provided to the

client system for participation. Thus, the user may be provided information as to how the user's incentives would increase or change depending upon other computer systems or upgraded capabilities the user could acquire. This incentive value increase may also be tied to upgrades to particular vendor's devices. For example, if the user's device is a computer system having an ABC microprocessor, the block 276 may provide the user information as to increased incentive values based upon an upgrade to a more powerful ABC microprocessor. Similarly, if the user's device is a computer system obtained from ABC, the block 276 may provide the user information as to increased incentive values based upon an upgrade to a more powerful ABC computer system.

FIG. 2D is an example user interface 276 for a client system agent, including incentive advertising, according to the present invention. In the example shown, interface 276 is a window 230 that may be displayed on a distributed device, for example, a computer system. This window 230 displays the desired information for the agent client manager. As indicated above, this agent client manager is initially downloaded from the server systems 104 and thereafter may be updated at various times when the client system is communicating with the server systems. The interface 276, as shown, includes interface tabs 221, 222, 224, 226, 228, 244, 246 and 248. These interface tabs may be selected through the user of a pointing device or keyboard attached, for example, to a computer system graphically displaying the window 230. It is noted that the interface tabs 221, 222, 224, 226, 228, 244, 246 and 248 are only examples, and the number, arrangement and content of tabs may be modified as desired. In addition, the example user interface 276 depicted in FIG. 2D is only an example and may be modified as desired.

In FIG. 2D, the processor values interface tab 224 is the one currently selected by the user. This tab 224 (Processor Values) includes example information that may be displayed to the user. Assuming that a workload is being processed by the agent client manager, the user may select the button 242 (Show My Incentive Values) to show the user's current incentive values associated with the workload being performed. The personal incentive values chart 232 (My Personal Incentive Values) may then be displayed to the user. As shown, the incentive values are provided in a relative scale from 1 to 10. The key designation 240 represents the incentives associated with the users current central processing unit (CPU) or microprocessor.

As indicated above, this incentive information may also be tied to the specific vendor of the user's CPU, for example, ABC Company's CPU. Thus, as shown, the key designation 240 (My current

processor) and the corresponding bar graph portion 236 represent incentives for the user's current CPU (e.g., a 166 MHz processor). The key designation 238 represents the incentives that the user is projected to have if the user were to upgrade the CPU. Again, this upgrade incentive information may be tied to the specific vendor of the user's CPU or to any other vendor, if desired. Thus, as shown, the key designation 238 (NEW ABC 1GHz processor!) and the corresponding bar graph portion 234 represent incentives for an upgrade to a new ABC CPU (e.g., a new ABC 1 GHz processor). In this manner, a user may be provided an incentive to increase the capabilities of the distributed device, and a vendor may be provided advertising so that the user is also directed to a particular upgrade.

Looking further to FIG. 2D, other similar incentive related information tabs may be provided for any desired capability of the distributed device. For example, tab 246 (Memory Values) represents information that may be provided for the memory capabilities of the distributed device. Tab 222 (Graphics Values) represents information that may be provided for the graphics capabilities of the distributed device. Tab 226 (Communications Values) represents information that may be provided for the communication capabilities of the distributed device. Tab 228 (Storage Values) represents information that may be provided for the storage capabilities of the distributed device. Tab 248 (System Values) represents information that may be provided for the system capabilities as a whole for the distributed device.

In addition to these incentive related information tabs, other tabs may be included to provide information and control for any desired features of the agent client manager. For example, the tab 244 (Current: Prime Search) represents information that may be displayed to the user about the current workload being performed by the agent client manager, for example, a search for large prime numbers. The tab 221 (Settings) represents information that may be displayed to the user about various settings for the client agent manager. In particular, the tab 221 may provide the user the ability to control any desired aspect of the operation of the agent client manager. For example, the user may be able to select a portion of the capabilities that may be utilized (e.g., a maximum of 20% of the system memory), the types of workloads that may be performed (e.g., only scientific research projects), the times when the agent may utilize system resources (e.g., only between 12 to 6 am, or only when the system is idle), or any other desired operational feature. It is noted that in addition to upgrade incentive information indicated above, the user may also be provided information as to how incentives would increase if the user allocated or changed the settings for the agent client manager.

0334785-01301
50
5
10
This user selection of operational features allows for workloads to be scheduled or balanced based upon user input and desires. These user vectors, as indicated above, would allow users to dedicate their device capabilities to specific research projects (cancer, Parkinson's disease, Internet, genetics, space science, etc.), to specific non-profit or for profit organizations (Greenpeace, Celera, etc.), educational institutions (University of Texas), a specific group of like minded users, or any other entity or endeavor. This affiliation selection allows the distributed processing system to automatically include a user's device capabilities in a pool dedicated to the chosen affiliation. Additionally, a user could choose to mix various percentages and allocations of device capabilities among multiple affiliations. It is noted that the user need not make any affiliation selection and need not allocate 100 percent of device capabilities. Rather, only a portion of the device capabilities may be allocated to a particular affiliation, leaving the remainder non-allocated and not affiliated. The capability allocation may also be a system-wide (i.e., course) allocation, such as some desired percent of overall device capabilities. The capabilities allocation may also be subsystem specific (i.e., fine) allocation, such as allocation of particular subsystem capabilities to particular affiliations.

Now looking to FIG. 3A, the server systems 104 may be one or more computer systems that operate to identify client system capabilities, organize workloads, and utilize client systems to accomplish a desired task. The server systems 104 includes a control system 304 a workload database 308, and a sweepstakes system 306, as discussed more below. The workload database 308 stores any desired project task, which may be broken up into discrete workload tasks WL1, WL2 ... WLN, as represented by elements 336, 338 ... 340. The workload database may also store one or more benchmark workloads (BWL) 335 that may be utilized to determine client system capabilities in response to a standard workload. Through line 312, the workload database 308 communicates with control system 304. Control system 304, for example, receives original workload 322 and transfers it to the interface 320 through line 326. The interface 320 then transfers the workload 322 to the network 102 through line 114. This workload 322 is ultimately received as workload 204 by client system 108, 110 or 112, as shown in FIG. 2A. The result 324 is ultimately received by the control system 304 through interface 320 and line 328.

30
In allocating workloads, the control system 304 may consider the capabilities of the client systems 108, 110 and 112 to which the control system 304 is sending workloads. For example, if client 108 has more processing power than client 110, the control system 304 may allocate and send more difficult or larger workloads. Thus, client 108 may receive WL1 336 and WL2 338, while client 110

would only receive WL3. Alternatively, the workload database 308 could be organized with differing levels of processing power or capability requirements for each workload. In this way, WL1 336 may represent a greater processing or system capability requirement than WL2 338. It should be noted that workload may be a processing task, a data storage task, or tied to any other of a variety of capabilities that may be utilized on the client systems 108, 110 ... 112.

As indicated above, to encourage owners or users of client systems to allow their system capabilities to be utilized by control system 304, an incentive system may be utilized. This incentive system may be designed as desired. Incentives may be provided to the user or owner of the clients systems when the client system is signed-up to participate in the distributed processing system, when the client system completes a workload for the distributed processing system, or any other time during the process. In addition, incentives may be based upon the capabilities of the client systems, based upon a benchmark workload that provides a standardized assessment of the capabilities of the client systems, or based upon any other desired criteria.

One example use of a benchmark workload is to use the benchmark workload to determine incentive values. For example, the server systems 104 may be designed to send out a standard benchmark workload once an hour to each client system 108, 110 ... 112. If a client system is not available at that time for any reason, the workload would not be completed by the client system, and there would be no incentive value generated for that client system. In this example, the benchmark workload may be a timed work-set that would exercise each subsystem with capabilities within the client system that was desired to be measured. A more capable client system would then generate greater incentive values from executing the benchmark workload, as compared to a less capable client system. These incentive values may be utilized as desired to determine what the client system should get in return for its efforts. For example, if the incentive were a sweepstakes as discussed further below, the number of entries in the sweepstakes may be tied to the system's performance of the benchmark workload. Thus, the faster or better the client system performs the benchmark workload, the more entries the client system would receive.

In the embodiment shown in FIG. 3A, the server systems 104 includes a sweepstakes system 306 that functions with control system 304 to provide incentives for the users or owners of client systems 108, 110 and 112 to allow their system capabilities to be used by the server systems 104. The control system 304 may determine a sweepstakes entry value 302 that is sent along line 310 to the sweepstakes

system 306. The sweepstakes system 306 may then receive sweepstakes entry 332 and provide it to the sweepstakes engine 330 through line 334. The sweepstakes engine 330 may process the entries and determine a winner, when desired. In the embodiment shown, therefore, entries to the sweepstakes may be generated each time a unit of work is accomplished by one or more of the subsystems within a client system 108, 110 or 112 via an agent installed on the device for the purposes of managing and completing units of work. The total entries for any period of time would, therefore, be dynamic depending on how many are received. Odds of winning would then be determined by the total number of entries received and the total number of entries contributable to any given entrant.

FIG. 3B is another example block diagram of a distributed processing system 300 including server systems 104, customer systems 152, client systems 134 and out-sourced host systems 340, according to the present invention. The server systems 104 may include an analytic subsystem 346, a results/workload production subsystem 344, a project pre-processing subsystem 342, a client agent subsystem 343, and an incentive advertising subsystem 345. The incentive advertising subsystem 345 may operate to provide advertising information, for example, the upgrade incentive information as discussed with respect to FIG. 2D. The client agent subsystem 343 may operate to download an agent to the client systems 134 and to update this agent at times when the server systems 104 are communicating with the client systems 134.

The customer systems 152, which represent customers that have projects that they desired to be processed by the distributed processing system, may be connected to the project pre-processing subsystem 342 to provide projects to the server systems 104. These projects are processed by the project pre-processing subsystem 342 and passed to the results/workloads production subsystem 344, which produces and sends out workloads 130 and receives back results 132. The analytic subsystem 346 then takes the results and processes them as desired. Completed project information may then be provided from the analytic system 346 to the customer systems 152. In this manner, the projects of the customer systems 152 may be processed and project results reported by the distributed processing system of the present invention.

Also, as shown, the workloads 130 and the results 132, or other tasks of the server systems 104, may be processed and handled by out-sourced host systems 340, if desired. Thus, some or all of the workloads 130 may be sent first to out-sourced host systems 340. Out-sourced host systems 340 then send workloads 130A to the client systems 134 and receive back results 132A. The out-sourced host

systems 340 then send the results 132 back to the server systems 104. It is noted that this out-sourcing of server system tasks may be implemented as desired for any given task that the server systems 104 may have. It is further noted that, if desired, the server systems 104 may perform all of the desired functions of the server systems 104 so that no out-sourced host systems 340 would be used.

5

FIG. 3C is a block diagram for one embodiment of a server system processor 350, according to the present invention. An agent abstraction layer 360 may send workloads 130 and receive results 132. The security subsystem 354 may interact with the agent abstraction layer 360 and provide information to a data parser 352 and an application programming interface (APIs) block 356. The APIs block 356, the data parser 352 and a workload manager 358 may interact to accomplish the desired tasks for the server system processor 350. It is noted that for this embodiment, the API protocol could be controlled and provided to other host systems.

FIG. 3D is an alternative block diagram for a server system processor 350, according to the present invention. In this embodiment, the APIs block 356 and the agent abstraction layer 360 are not present. The data parser 352, the workload manager 358 and the security subsystem 354 interact to provide the desired server system tasks. It is noted that for this embodiment, the security subsystem is controlled and utilized for communicating with client systems.

FIG. 4 is a functional block diagram for a sweepstakes operation 400 by the system server 104 according to the present invention. In block 402, the server systems 104 may sign-up client systems in "accept clients" block 402. Following line 418, the server systems 104 identifies the capabilities of the client's computer and processing systems in the "determine client system capabilities" block 404. Control passes along line 420 to the "distribute workloads to client systems" block 406, where the server systems 104 allocates workloads to each client system 108, 110 and 112. This workload may also be an benchmark workload, as indicated above, that acts as an entry workload to determine the entries or entry values for the client system. As also indicated above, in distributing the workloads in block 406, the server system 104 may take into consideration the capabilities of the client systems to which workloads are being distributed. The client systems 108, 110 and 112 then operate to complete the workloads allocated to them. Along line 423, the server system 104 receives back workload results in "receive workload results" block 408.

At this point, control passes along line 424 to the "determine sweepstakes entries" block 410. In this block 410, the server system 104 determines the entry value for the workload completed or for a standard benchmark or entry workload completed. This entry value may be weighted upon a variety of factors including factors such as the amount of work completed, the difficulty level of the processing required, and the accuracy of the results. It is noted that any desired weighting may be utilized. Thus, it is understood that a wide variety of considerations may be utilized to determine the entry value weighting for the sweepstakes.

Although the weighting determination is shown in block 410 in FIG. 4, the entry value may also be determined, in whole or in part, when a client system signs on to the distributed processing distributed system of the present invention. For example, if a client system has state-of-the-art CPU, video processor, DSP engine, memory, and large amounts of free disk storage space, a high entry value may be allocated to this client system up-front. In contrast, a client system that has a slow CPU, a weak video processor, no DSP engine, little memory, and little free disk storage space may be allocated a small entry value. In this way, the owners or users of the client systems may be provided immediate feedback as to the potential sweepstakes entry value of their computer systems, devices and system capabilities.

It is further noted that the entry value may take any desired form and may be, for example, a multiplier that will be used for each unit of workload completed. In this way, the owner or user will readily be cognizant that a state-of-the-art system will yield a high multiplier, where as an older system, system capability or device will yield a low multiplier. Such feedback, whether communicated to the owner or user immediately upon signing up or upon completion of each workload, will create an incentive for owners and/or users to acquire state-of-the-art systems, thereby further increasing the potential processing power of the distributed processing system of the present invention.

In addition, different workload projects may be designated with different entry values, as well. For example, some workload projects may require particular hardware or software processing systems within a client system or device. Thus, the number of client systems that are capable of performing the task would be limited. To further encourage participation by those owners or users with capable systems, the entry value for taking on particular workloads and/or systems with the desired features may be allocated higher entry values.

Referring back to FIG. 4, control passes along line 426 to the "process entries" block 412. In this block 412, the sweepstakes entries are processed and stored as desired. Following line 428, "end of entry period" decision block 414 represents a determination of whether the time for getting entries into the sweepstakes has ended. If not, the control continues to line 430 and back to blocks 402, 404 and/or 406, depending upon what is desired. Once the entry period has ended, control flows along line 432 to "determine winners" block 416. The server system 104 then identifies from among the entries, who the winning client system or systems will be.

The entry period may be any desired time frame and may include multiple overlapping time frames, as desired. For example, winners may be determined daily for entries each day, monthly for entries within a month, and/or yearly for entries within one year. In addition, special entry periods may be generated, if desired, for example where a particularly important workload project had a short time frame in which it needed to be completed.

FIGS. 1, 2A-C, 3A-D, and 4 are directed to example embodiments for a distributed processing system according to the present invention, including a sweepstakes reward or incentive feature, as shown in the embodiments of FIG. 3A and FIG. 4.

FIGS. 6A and 6B further describe a capabilities scheduling feature, in which the server systems 104 may identify and consider any of a variety of client system capability vectors in determining how to organize, allocate and manage workloads and projects. FIGS. 5A and 5B describe a distributed processing system and workload project that accomplishes network site indexing. FIGS. 7A and 7B describe a distributed processing system and a workload project that accomplishes network site testing, such as quality of service (QoS) testing and load testing. And FIG. 8 describes a distributed processing system, preferably with respect to a corporate intranet, that accomplishes distributed data back-up.

FIG. 9 is an alternative representation for the interconnection fabric for a distributed processing system environment and describes idle client system identification and shared component client systems. FIG. 10 describes a client system agent installed on a client system. FIGS. 11A and 11B further describe machine generated sweepstakes entries. FIGS. 12A and 12B describe client capability selection features. FIGS. 13A and 13B describe data conversion services. FIG. 14A describes a distributed processing system that provides data transmission caching. FIG. 14B describes a distributed processing system that

provides data sharing and file distribution functions. And FIG. 15 describes an alternative representation for a distributed processing system, according to the present invention.

Looking now to FIG. 5A, block diagram is depicted of a distributed processing system 550 for a network site indexing application, according to the present invention. As stated above with respect to FIG. 1A, the network 102 may be a wide variety of networks. For this network site indexing application, the network 102 may preferably be the Internet having a multitude of network sites 552 ... 554. Each network site 552 ... 554 may have a variety of different content types that may be indexed, ranging from complex sites to relatively simple sites. For example, network site 552 includes text 570A, images 570B, audio streams 570C, video streams 570D, files 570E and other content 570F. Network site 554 is less complex and includes text 572A, images 572B, and other content 572C. Both network sites 552 and 554 are connected to the network 102 through communication lines 558 and 556, respectively.

As discussed above, the server systems 104 manage workloads for the client systems 108, 110 ... 112. The client systems 108, 110 ... 112 process these workloads and produce indexing results. The resulting index may be stored at a centrally managed site, such as central index storage block 560, or may itself be distributed over the possibly millions of indexing clients 108, 110 ... 112, as shown by remote index storage blocks 562, 564 ... 566. If remote index storage is utilized, a master database content index may be stored locally, for example, in the central index storage block 560. This content index may then direct relevant searches to the distributed massively parallel engine for search queries.

Referring now to FIG. 5B, a functional block diagram is shown for a network site indexing operation 500 according to the present invention. As described in FIG. 1A with respect to other systems 106, there may be any number of computer and processing systems connected to the network 102. Any one of these others systems 106 may publish information on the network 102 for access by any other system connected to the network 102. This information to be indexed may take a wide variety of forms, including, for example, text, images, audio streams, video streams, databases, spreadsheets, PDF files, Shockwave data, Flash data, applications, data files, chat streams, or any other information, data or data streams that may be accessible on a network site. The distributed processing system of the present invention may have as a workload the task of indexing this potentially massive amount of information.

For example, where the network 102 is the Internet or a large intranet, a large amount of processing power and time is needed to create an accurate, complete and up-to-date index of the

information. The Internet uses an IP (Internet Protocol) address protocol to direct traffic around the Internet. The IP address is the address of a computer attached to a TCP/IP (Transmission Control Protocol/Internet Protocol) network. Every system on the network must have a unique IP address. IP addresses are typically written as four sets of numbers separated by periods. The TCP/IP packet uses 32 bits to contain the IP address, which is made up of a network and host address (NETID and HOSTID). The more bits used for network address, the fewer remain for hosts. Web pages within a particular web site with a unique address may be addressed through URLs (Uniform Resource Locator) associated with that web site. In short, there is a limited, but very large, number of possible IP addresses for uniquely identifiable Internet sites that may be accessed and analyzed to generate an index of Internet sites and web pages via URLs.

The operation diagram of FIG. 5B starts with the "clients receive indexing workloads" block 502. In this block, the system server 104 provides the clients systems 108, 110 ... 112 with a workload task to index a portion of the information accessible on the network 102. For example, with the Internet, each workload may be single IP address or groups of URLs or, in some cases, large data types contained on single sites or pages. Following line 514, the "clients interact with other systems" block 504 represents the operation of the agent installed on the client systems 108, 110 ... 112 to access the network sites, according to the assigned workload, and index the information accessible on that site. This indexing may include all types of information accessible on that site, including text, audio, image, video, etc.

Next, following lines 516 and 518, the client systems 108, 110 and 112 complete the workload tasks, get the results ready for transmission, and sends those results back to the system server 104 in "clients complete workload" block 506 and "indexing results sent to server system" block 508. Control passes along line 520 to "index compiled for use" block 510 where the server system formats and/or compiles the results for use. For example, the index results may be utilized for accurate, complete and up-to-date search information for the network 102. As indicated with respect to FIG. 5A, the resulting index may be stored remotely or locally following line 522. Thus, element 524 represents remote storage of the index, and element 526 represents central storage of the index. It is noted that the index may also be stored with a mixture of central and remote storage, as desired. In addition, as indicated above, a directory or summary index for the resulting index may be generated and stored centrally, if desired. It is further noted that the summary index may be stored in any other desired fashion, for example, it may be distributed and stored on a number of client systems.

FIG. 6A is a block diagram for a server system 104 according to the present invention, including a control system 304, a workload database 308, and a database of capability vectors 620. The workload database 308 includes a variety of sets of workload projects WL1, WL2 ... WLN. For each workload project, there may be multiple workload units. For example, workload project WL1 includes workload units WL11, WL12 ... WL1N, as represented by elements 640, 642 ... 644, respectively. Similarly, workload project WL2 includes workload units WL21, WL22 ... WL2N, as represented by elements 646, 648 ... 650, respectively workload project WLN includes workload units WLN1, WLN2 ... WLNn, as represented by elements 652, 654 ... 656, respectively.

It may be expected that different workload projects WL1, WL2 ... WLN within the workload database 308 may require widely varying processing requirements. Thus, in order to better direct resources to workload projects, the server system may access various system vectors when a client system signs up to provide processing time and other system or device capabilities to the server system. This capability scheduling helps facilitate project operation and completion. In this respect, the capability vector database 620 keeps track of any desired feature of client systems or devices in capability vectors CBV1, CBV2 ... CBVn, represented by elements 628, 630 ... 632, respectively. These capability vectors may then be utilized by the control system 304 through line 626 to capability balance workloads.

This capability scheduling according to the present invention, therefore, allows for the efficient management of the distributed processing system of the present invention. This capability scheduling and distribution will help maximize throughput, deliver timely responses for sensitive workloads, calculate redundancy factors when necessary, and in general, help optimize the distributed processing computing system of the present invention. The following TABLE 1 provides lists of capability vectors or factors that may be utilized. It is noted that this list is an example list, and any number of vectors or factors may be identified and utilized, as desired.

TABLE 1 -- Example Client Capability Vectors or Factors

1. BIOS Support:	a. BIOS Type (brand)
b. ACPI	c. S1, S2, S3, and S4 sleep/wake states
d. D1, D2 and D3 ACPI device states	e. Remote Wake Up Via Modem
f. Remote Wake Up Via Network	g. CPU Clock control
h. Thermal Management control	i. Docked/Undocked state control
j. APM 1.2 support	k. Hotkey support
l. Resume on Alarm, Modem Ring and LAN	m. Password Protected Resume from Suspend
n. Full-On power mode	o. APM / Hardware Doze mode
p. Stand-by mode	q. Suspend to DRAM mode
r. Video Logic Power Down	s. HDD, FDD and FDC Power Down
t. Sound Chip Power Down	u. Super I/O Chip Power Down
2. CPU Support:	a. CPU Type (brand)
b. MMX instruction set	c. SIMD instruction set
d. WNI instruction set	e. 3DNow instruction set
f. Other processor dependent instruction set(s)	g. Raw integer performance
h. Raw FPU performance	i. CPU L1 data cache size
j. CPU L1 instruction cache size	k. CPU L2 cache size
l. CPU speed (MHz/GHz...)	m. System bus (MHz/GHz...) speed supported
n. Processor Serial Number	o. CUID
3. Graphic Support	a. Graphics type (brand)
b. # of graphics engines	c. Memory capacity
d. OpenGL support	e. Direct3D/DirectX support
f. Color depth supported	g. MPEG 1/II decode assist
h. MPEG1/II encode assist	i. OS support
j. Rendering type(s) supported	k. Single-Pass Multitexturing support
l. True Color Rendering	m. Triangle Setup Engine

n. Texture Cache	o. Bilinear/Trilinear Filtering
p. Anti-aliasing support	q. Texture Compositing
r. Texture Decompression	s. Perspectively Correct Texture Mapping
t. Mip-Mapping	u. Z-buffering and Double-buffering support
v. Bump mapping	w. Fog effects
x. Texture lighting	y. Video texture support
z. Reflection support	aa. Shadows support
4. Storage Support	a. Storage Type (brand)
b. Storage Type (fixed, removable, etc.)	c. Total storage capacity
d. Free space	e. Throughput speed
f. Seek time	g. User dedicated space for current workload
h. SMART capable	
5. System	a. System Type (brand)
b. System form factor (desktop, portable, workstation, server, etc.)	
6. Communications Support	a. Type of Connection (brand of ISP)
b. Type of Connection Device (brand of hardware)	c. Hardware device capabilities
d. Speed of connection	e. Latency of connection
f. Round trip packet time of connection	g. Number of hops on connection type
h. Automatic connection support (yes/no)	i. Dial-up only (yes/no)
j. Broadband type (brand)	k. Broadband connection type (DSL/Sat./Cable/T1/Intranet/etc.)
7. Memory	a. Type of memory error correction (none, ECC, etc.)

b. Type of memory supported (EDO, SDRAM, RDRAM, etc.)	c. Amount of total memory
d. Amount of free memory	e. Current virtual memory size
f. Total available virtual memory size	
8. Operating System	a. Type of operating system (brand)
b. Version of operating system	c. Health of operating system
9. System application software	a. Type of software loaded and/or operating on system
b. Version of software	c. Software features enabled/disabled
d. Health of software operation	

FIG. 6B is a functional block diagram for capabilities determination and scheduling operation 600 for workloads in a distributed processing system according to the present invention. Initially, various vectors are identified for which capability information is desired in the "identify client system capability vectors" block 602. Following line 612, the server systems 104 then capability balances workloads among client systems 108, 110 and 112 based upon the capability vectors in the "capability scheduling workloads based on vectors" block 604. Then the capabilities scheduled workloads are sent to the client systems for processing in the "send capability scheduled workloads" block 606.

This capability scheduling and management based upon system related vectors allows for efficient use of resources. For example, utilizing the operating system or software vectors, workloads may be scheduled or managed so that desired hardware and software configurations are utilized. This scheduling based upon software vectors may be helpful because different software versions often have different capabilities. For example, various additional features and services are included in MICROSOFT WINDOWS '98 as compared with MICROSOFT WINDOWS '95. Any one of these additional functions or services may be desired for a particular workload that is to be hosted on a particular client system device. Software and operating system vectors also allow for customers to select a wide variety of software configurations on which the customers may desire a particular workload to be run. These varied software configurations may be helpful, for example, where software testing is

desired. Thus, the distributed processing system of the present invention may be utilized to test new software, data files, Java programs or other software on a wide variety of hardware platforms, software platforms and software versions. For example, a Java program may be tested on a wide proliferation of JREs (Java Runtime Engines) associated with a wide variety of operating systems and machine types, such as personal computers, handheld devices, etc.

From the customer system perspective, the capability management and the capability database, as well as information concerning users of the distributed devices, provide a vehicle through which a customer may select particular hardware, software, user or other configurations, in which the customer is interested. In other words, utilizing the massively parallel distributed processing system of the present invention, a wide variety of selectable distributed device attributes, including information concerning users of the distributed devices, may be provided to a customer with respect to any project, advertising, or other information or activity a customer may have to be processed or distributed.

For example, a customer may desire to advertise certain goods or services to distributed devices that have certain attributes, such as particular device capabilities or particular characteristics for users of those distributed devices. Based upon selected attributes, a set of distributed devices may be identified for receipt of advertising messages. These messages may be displayed to a user of the distributed device through a browser, the client agent, or any other software that is executing either directly or remotely on the distributed device. Thus, a customer may target particular machine specific device or user attributes for particular advertising messages. For example, users with particular demographic information may be targeted for particular advertisements. As another example, the client agent running on client systems that are personal computers may determine systems that are suffering from numerous page faults (i.e., through tracking operating system health features such as the number of page faults). High numbers of page faults are an indication of low memory. Thus, memory manufacturers could target such systems for memory upgrade banners or advertisements.

Still further, if a customer desires to run a workload on specific device types, specific hardware platforms, specific operating systems, etc., the customer may then select these features and thereby select a subset of the distributed client systems on which to send a project workload. Such a project would be, for example, if a customer wanted to run a first set of simulations on personal computers with AMD ATHLON microprocessors and a second set of simulations on personal computers with INTEL PENTIUM III microprocessors. Alternatively, if a customer is not interested in particular configurations

for the project, the customer may simply request any random number of distributed devices to process its project workloads.

Customer pricing levels for distributed processing may then be tied, if desired, to the level of specificity desired by a particular customer. For example, a customer may contract for a block of 10,000 random distributed devices for a base amount. The customer may later decide for an additional or different price to utilize one or more capability vectors in selecting a number of devices for processing its project. Further, a customer may request that a number of distributed devices be dedicated solely to processing its project workloads. In short, once device attributes, including device capabilities and user information, are identified, according to the present invention, any number of customer offerings may be made based upon the device attributes for the connected distributed devices. It is noted that to facilitate use of the device capabilities and user information, capability vectors and user information may be stored and organized in a database, as discussed above.

Referring now to FIG. 12A, a block diagram depicts a distributed processing system 1200 that allows customers to select client system attributes, such as device capabilities and user characteristics, according to the present invention. In this embodiment, the network 102 is depicted as the Internet to which server systems 104, customer 152A, customer 152B, and client systems 1202A, 1202B ... 1202C are connected. These systems are connected through communication links 114, 119A, 119B, 1204A, 1204B ... 1204C, respectively. As noted above, these communication links may include any of a wide variety of devices and/or communication techniques for allowing a system to interface with other connected systems.

As shown in FIG. 12A, and as discussed above, the customers 152A and 152B may desire to send information or projects, such as advertisements (ADV) 1206A and 1206B and/or projects (PROJ) 1208A and 1208B, to groups of client systems that have particular or selected capabilities. The number of different groups of client systems is as varied as the capability and user data available for those client systems. The client systems 1202A represent client systems that include a first set (Set 1) of desired attributes. The client systems 1202B represent client systems that include a second set (Set 2) of desired attributes. And the client systems 1202C represent client systems that include a Nth set (Set N) of desired attributes. Once attributes are selected, the client systems with those attributes may be accessed as desired by customers 152A and 152B. For example, customer 152A may send its advertisement to client systems 1202B. Customer 152B may send its advertisement to client systems 1202A. The project

1208A from customer 152A may be processed by client systems 1202C. And the project 1208B from customer 152B may be processed by client systems 1202B. It is noted, therefore, that any combination of desired attributes, such as device capabilities and user characteristics, may be identified and utilized to satisfy customer objectives, whether those objectives be advertising, project processing, or some other desired objective.

FIG. 12B is a block flow diagram for client system attribute selection, according to the present invention. In the embodiment shown, process 1250 begins with the customer selecting desired attributes in block 1252. Next, client systems with selected attributes are accessed in block 1254. And, then in block 1256, the customer objective, such as advertising or project, is processed by the client system. Control of this process 1250 may be provided by the server systems 104, if desired, such that the customer interfaces with the server systems 104 to select device attributes and then the servers systems 104 access the client systems. Alternatively, the server systems 104 may simply provide the customer with a list of contact information (e.g., IP addresses) for the client systems, so that the customer may directly access the client system, for example, in providing advertisements to the users of the client systems. It is further noted that other control techniques may also be used to identify and access client systems with particular desired device capabilities, user characteristics, or other device attributes, according to the client system attribute selection method of the present invention.

FIG. 7A is a block diagram for a distributed processing system 100 according to the present invention, including example network sites 106A and 106B on which site testing is to be conducted, such as load testing and/or quality-of-service (QoS) testing. FIG. 7A is similar to FIG. 1A except that other systems 106 in FIG. 1A has been represented in the embodiment of FIG. 7A with network sites 106A and 106B. Communication line 116A between the network 102 and the network site 106A represents a interaction by one client system 108, 110 and 112. Communication lines 116B, 116C and 116D represent interactions by more than one client system 108, 110 and 112.

Site testing is typically desired to determine how a site or connected service performs under any desired set of test circumstances. With the distributed processing system of the present invention, site performance testing may be conducted using any number of real client systems 108, 110 and 112, rather than simulated activity that is currently available. Several tests that are commonly desired are site load tests and quality of service (QoS) tests. Quality of service (QoS) testing refers to testing a user's experience accessing a network site under normal usability situations. Load testing refers to testing what

a particular network site's infrastructure can handle in user interactions. An extreme version of load testing is a denial-of-service attack, where a system or group of systems intentionally attempt to overload and shut-down a network site. Advantageously, the current invention will have actual systems testing network web sites, as opposed to simulated tests for which others in the industry are capable.

5

Network site 106B and the multiple interactions represented by communication lines 116B, 116C and 116D are intended to represent a load testing environment. Network site 106A and the single interaction 116A is indicative of a user interaction or QoS testing environment. It is noted that load testing, QoS testing and any other site testing may be conducted with any number of interactions from client systems desired, and the timing of those interactions may be manipulated and controlled to achieve any desired testing parameters. It is further noted that periodically new load and breakdown statistics will be provided for capacity planning.

10

FIG. 7B is a functional block diagram for a site-testing operation 700 according to the present invention. Initially, client systems 108, 110 and 112 receive workloads that identify testing procedures and parameters in the "clients receive testing workload" block 702. Following line 714, the client systems 108, 110 and 112 access the site being tested and perform the testing in block "clients interact with other systems" block 704. Next, following lines 716 and 718, the client systems 108, 110 and 112 complete the site testing workload tasks, get the results ready for transmission, and send those results back to the system server 104 in "clients complete testing workload" block 706 and "site testing results sent to server system" block 708. Control passes along line 720 to "site testing results compiled for use" block 710 where the server system formats and/or compiles the results for use by the network site. For example, the site testing results may be utilized determining modifications that need to be made to the network site to handle peak volume activities.

25

FIG. 8 is a block diagram for a distributed processing system 800 for a data back-up system application, according to the present invention. As stated above with respect to FIG. 1A, the network 102 may be a wide variety of networks, including an intranet network. Intranet networks, such as internal networks set up by corporations, are particularly suited for this application because the systems holding the data being backed-up would be owned by the same entity owning other systems with excess data storage capabilities. In this way, security would not be as great of an issue and the client system types could be better controlled. It is noted, however, that this data back-up application would be equally applicable to other networks, such as for computer systems connected through the Internet.

30

0034785-01301
FOET-00-58743360

Referring back to FIG. 8, client systems 108, 110 ... 112 are shown each having a back-up data blocks 804, 806 ... 808. Customer systems 152 is shown as having data 802, which is desired to be backed-up with the distributed back-up system 800. The server systems 104 manage the flow of data from the data 802 and the client systems that have extra storage space represented by back-up data blocks 804, 806 ... 808. In operation, the server systems 104 identifies client system storage capabilities. With this information, the server systems 104 can receive data for back-up from any system on the network 102. It is noted, and as indicated with respect to FIG. 1A, the client systems 108, 110 ... 112 and the customer systems 152 may communicate directly with each other in peer-to-peer type communications.

The server systems 104 may also manage the storage and transfer of data so that the data will be readily retrievable once backed-up and stored on the client systems 108, 110 ... 112. If desired, an summary index or directory of the backed-up data may be stored centrally on the server systems 104, or may be stored remotely on the client systems 108, 110 ... 112. It is also noted that the server systems 104 may also distribute data back-up workloads so that each portion of the data 802 is stored redundantly on at least two of the client systems 108, 110 ... 112. This redundancy provides added security should any one or more client systems suddenly cease to be operational.

Looking now to FIG. 9, a block diagram is depicted of an alternative representation of an interconnection fabric for a distributed processing system environment 100, according to the present invention. In this diagram and as described above, the network environment may be the Internet, an internal company intranet, a local area network (LAN), a wide area network (WAN), a wireless network, a home network, or any other system that connects together multiple systems and devices. In addition, the server systems and clients systems may be interconnected by a variety of possible connection interfaces, for example, Ethernet connections, wireless connections, ISDN connections, DSL connections, modem dial-up connections, cable modem connections, direct T1 or T3 connections, fiber optic connections, routers, portal computers, as well as any other network or communication connection. It is noted, therefore, as discussed with respect to other embodiments such as the embodiment of FIG. 1A, that systems may be coupled into an interconnected fabric in any of a variety of ways and communications can potentially occur directly or indirectly between any of the systems coupled into the fabric, as would be understood by those of skill in the art.

Within this environment, as depicted in FIG. 9, server systems 104 are interconnected with any number of client systems, for example, client systems 108A, 108B, 108C, 108D, 108E, 108F, 108G, 108H, 108I, 108J, 108K and 108L. In addition, these client systems may also include idle client systems 902A, 902 B, and 902C, as discussed further below. Furthermore, these client systems may include client system 904A with a component A, client system 904B with a component B, and client system 904C with a component C. It is also noted that the interconnection fabric may include any number of devices that are not client systems, in that they themselves are not providing components or processing capabilities for the distributed processing system of the present invention. Nevertheless, these devices may be considered part of the system because they may relay, interpret, process or otherwise transmit or receive information from or to client systems that are part of the distributed processing system.

Aggregation of component level resources, according to the present invention, will now be discussed. As described above, the capabilities of client systems are determined for purposes of allocating, scheduling and managing distributed processing workloads. In other words, each of the client systems may be made up of many individual subsystems with various capabilities. In some cases, it may occur that particular components on different machines may provide added value if combined or aggregated. Thus, utilizing subsystem or component level resources from a heterogeneous group of devices may be the most efficient or otherwise advantageous way of taking advantage of these resources to complete various desired tasks.

Referring now more particularly to FIG. 9, the client systems 904A, 904B and 904C may have component A, component B and component C, respectively, that are better utilized in combination. For example, client system 904A may have a fast processor, a high-speed network connection, but little available storage space. Client system 904B may have large amounts of available free storage space but little processing power. Client system 904C may also have a fast processor, but relatively little available storage space. In this example, a workload that requires both a large storage capacity and a fast processor may be efficiently completed by dedicating component level resources to various parts of the workload from different machines. Thus, the workload may be managed by having client systems 904A and 904C processing data stored on and transmitted from client system 904B. Once clients systems 904A and 904C process data, this resulting data may then be transmitted back to client system 904B for aggregation and eventual transmission back to the server systems 104. The client system 904B, therefore, essentially acts as a server for a workload subset, sending out portions of a subset workload, receiving back the processed data, and aggregating the data to build a completed workload subset.

It is noted that any number of different components from different client systems may be aggregated, as desired. For example, for wireless devices, DSP processing and storage components could be aggregated with components from other client systems. For display devices, graphics rendering power could be aggregated. For relatively dumb machines, such as connected household appliances, vending machines, etc., slow-speed processing components could be aggregated. In short, an appropriate workload may include instructions to numerous client systems that will enable collaboration and aggregation of component level resources. Such instructions may include things, such as, where to receive input, where to send output, and ultimately which client systems return final results.

10

It is further noted that the control instructions may be de-centralized as well. In other words, as indicated above, client systems may communicate directly with each other, for example, in a peer-to-peer fashion. In this way, workload communications may occur directly between client systems, and workload control and management may occur through the client system agents located on client systems.

Still referring to FIG. 9, idle system determination will now be discussed. As stated above, client system capabilities are determined and utilized within the distributed processing system of the present invention. The more idle any particular client system, the more processing it is arguably able to accomplish, and the more incentives it is likely to receive. In other words, the client system capabilities may be utilized more often and more intensely if the client system is more idle. As such, it is advantageous to identify idle client systems and allocate them to more processor and time sensitive tasks. By identifying these idle client systems, resources available on the network at any given time may be more fully utilized, and otherwise idle resources may be utilized for highly intensive, real-time activities that would otherwise require dedicated devices. Examples of such real-time activities include data caching, indexing, etc. In FIG. 9, idle client systems are designated as 902A, 902B and 902C.

25

Identifying idle resources may be determined in any of a variety of ways. It is possible, for example, to simply look at whether a machine is not being used or has low processor utilization at any given time. This simple determination, however, may not yield an accurate picture of how idle a client system may or may not be over a given time period. More particularly, discovery methods may be implemented to identify the activity of a variety of client system components and subsystems. For example, subsystems may be monitored, such as network activity, device output activity, user input, processing activity, executing task monitoring, or mode of operation parameters (e.g., mobile or power

30

d. screen output (e.g., time elapsed since last screen output, paint event or pixel change)	e. network or communication packets sent or received (e.g., time elapsed since last network or communications activity)
f. storage device activity (e.g., time elapsed since last storage device activity, such as hard drives, flash memory cards, removable drives, CD drives, DVD drives, etc.)	g. processor, DSP, microcontroller, embedded device, or other processor activity (e.g., time elapsed since last processor activity)
h. processor, DSP, microcontroller, embedded device, or other processing device utilization (e.g., change in utilization levels)	i. tasks or processes executing (e.g., time elapsed since change in number of tasks or processes executing)
j. task or process device utilization (e.g., time since change in task or process device utilization)	k. any other device activity that could be used to classify if a device is idle

As a further example of the usefulness of this determination, reference is made back to FIG. 9. Server systems 104 may have, for example, a large, intensive task that it would like to place on these idle devices. After using a number of the vectors in TABLE 2 to determine the utilization level for client systems, the server systems 104 determines that client systems 902A, 902B and 902C are idle and capable of handling significant time sensitive processing tasks. For example, idle client systems 902A, 902B and 902C may be personal computers that can act as a local internet cache for other connected devices, such as some of the other client systems depicted in FIG. 9, that are interested in a data type that benefits from a local network cache. Thus, data or content may be transmitted from a remote network site to the idle machines 902A, 902B and 902C. These idle devices 902A, 902B and 902C may then re-transmit this same data or content to other connected devices also interested in the data or content.

One example for such network caching is Internet video or multimedia broadcast events that are desired to be viewed or received by a very large number of geographically close connected devices at about the same time. In order to meet the demand of these connected devices, web sites broadcasting an event have to be able to handle a huge increase in network traffic over a short period of time. By locally

5 caching the transmission to idle client systems, a web site can reduce the direct demand on its own resources. This is so because other connected devices may receive a re-transmitted broadcast, although delayed, from the idle client system. It is noted that according to the present invention idle client systems 902A, 902B and 902C may work independently or in combination. Even though idle client systems are suited for providing the caching function, it is also noted that that network caching may be accomplished using one or more client systems regardless of their respective levels of idleness.

10 FIG. 10 is a more detailed block diagram for a client system agent 270 installed on a client system, according to the present invention. This diagram includes a security subsystem 1010, a capabilities subsystem 1006, a workload processor 1004, a user interface 1002, and a project management and agent control subsystem 1008. The various components and subsystems may communicate with each other, for example, through lines 1012, 1014, 1016, 1018 and 1020. Externally, the client system agent 270 may communicate through its security subsystem 1010 with the other components within the client system and ultimately to other devices connected into the network fabric. It is noted that configuration of the client system agent and its operation, both internal and external, may be selected and designed, as desired.

As depicted, the capabilities subsystem 1006 includes an idle system monitor 1022, as described above, that monitors and analyzes user and device activities associated with the client system to determine the level of activity or idleness for the client system. The information determined by this idle system monitor 1022 may then be communicated externally, for example, through the security subsystem 1010 to the server systems 104. The server systems 104 may then store and analyze system idleness data from across the distributed processing system. This idleness data may become part of the capabilities database that is utilized to allocate and manage workloads and processing system resources.

25 Still referring to FIG. 10, the workload processor 1004 includes a machine entry generation subsystem 1024. As described above, the workload processor 1004 may send completed workloads back to server systems 104 to generate sweepstakes entries for the host client system. In this way, when the incentive is a sweepstakes, the client system may generate entries by completing workloads. The machine entry generation subsystem 1024 refers to this entry generation through workload completion. As discussed above, the workload processed to generate entries may be a project workload, an entry workload, or any other workload, as desired.

FIGS. 11A and 11B provide more detailed flow diagrams of process embodiments for machine generated sweepstakes entries through processing of entry workloads, according to the present invention.

Looking first to FIG. 11A, an entry workload process flow 1100 is depicted that provides machine generated sweepstakes entries. Process moves from start block 1102 to block 1104 in which entry workloads are loaded on client systems. Next, process flows to block 1106 which represents a periodic timer or other timing control for entry workload processing. After this timing control, the client system executes or processes the entry workload in block 1108. In block 1110, a sweepstakes entry is thereby generated and returned to the server system 104 based upon the completion of this entry workload. Process control then may proceed back to the periodic timing block 1106, where timing control determines when the entry workload is next processed. The completed workload represents the machine generated sweepstakes entry.

FIG. 11B is an alternative entry workload process flow 1150. The process flow 1150 is similar to the process flow 1100 except that the entry workload is sent to the client system each time it is to be run. Process starts in block 1102 and passes to the periodic timer block 1106, in which the process is controlled. For example, server systems 104 may determine when it is desirable for the client systems to receive and process an entry workload. In block 1104, the entry workload is sent to the client systems. As with FIG. 11A, the client systems then execute the entry workload in block 1108, and an entry is generated and returned to the remote server systems 104 in block 1110. The process then proceeds back to the periodic timer 1106 until it is determined that another entry workload should be processed. The primary difference between process 1150 and process 1100 is that process 1150 is depicting an entry workload that is transmitted to the client system each time it is to be run.

One example utilizing the process 1150 or the process 1100 is for server systems 104 to query the client systems for entry workload processing at regular time intervals. If a distributed device returns a completed entry workload back within a selected period of time from the distribution of the entry workload, the server system may conclude that the distributed device should receive an entry because the distributed device is providing resources to the distributed processing system. In this way, the server systems 104 may determine at regular intervals whether a given client system is working on project workloads for the distributed processing system. Alternatively, the client system agent may locally control the workload processing and may, for example, cause the client system to process and generate

entries at regular time intervals. It is noted that non-regular and varying time intervals may also be utilized and that combinations of remote and local control may also be utilized, as desired.

The timing of when a client system processes the entry workload, therefore, may be determined locally by the client system agent or remotely, for example, through commands sent by the server systems 104. In addition, periodic timing control may also be accomplished through various combinations of control routines residing locally and remotely. It is further noted that any number of different variations may be utilized to provide machine generated entries to a sweepstakes, according to the present invention. Thus, a client system may generate sweepstakes entries in any of a variety of ways and still have machine generated sweepstakes entries, according to the present invention.

FIGS. 13A and 13B describe a data conversion application 1300 for a massively parallel distributed network according the present invention. In particular, FIG. 13A is a block diagram of a distributed processing system that provides data conversion services, according to the present invention. And FIG. 13B is a block flow diagram for data conversion services within a distributed processing system, according to the present invention.

Converting file types, web pages, graphics images, etc., between device types can be a highly intensive processing task. Example devices that often need converted data are wireless devices, such as pagers and cell phones, that request Internet web page information from their respective device servers. The device server, instead of incurring the overhead of reformatting the requested data for the wireless devices, may instead distribute the requested page or data address, the device type information of the requesting device, and return address for the reformatted data. According to the present invention, the data conversion, translation or processing may be performed by a client system of the distributed processing system of the present invention. The resulting data may then be returned or provided to the original requesting device. In addition to data formatting for cell phones, language conversion, text translation and media translation services, or any other desired data conversion can also be hosted for a customer through the distributed processing system of the present invention.

It is noted that the data conversion operation contemplated by the present invention is not limited to any particular requesting device, any particular service provider, any particular type of data to be processed, any particular type of resulting processed data, or any particular data source. Thus, the data processed may include voice, text, application, image, source code, or any other data type or combination

of data types, and the resulting processed data may also include voice, text, application, image, or any other data type or combination of data types. According to the present invention, the distributed processing system is utilized to process any data that is desired by a requesting device and that must be converted or processed before being provided to the requesting device. For example, end-user devices
5 connected to the Internet, such as personal computers, may sign up for data conversion services through the server system so that the end-user device may request data conversion of any desired data, file, web site content, etc. Language translations and data formatting for connected wireless are just two examples of such applications for the present invention.

Looking now to the embodiment of FIG. 13A, the network 102 is depicted as the Internet, and the requesting device is one or more wireless devices 1306 connected to the Internet 102 through communication links 1308 and to the wireless device server systems 1304 through communication link 1309. The data to be converted, translated or otherwise processed is represented by block 1302 and may be, for example, content from an Internet web site that is connected to the Internet through communication link 1312. Also, as shown in FIG. 13A, a massively parallel distributed network (MPDN) server 104 is connected to the Internet 102 through communication link 114. The wireless device server systems 1304, or any other connected system that desires to off-load data conversion processing requirements (e.g., web site content servers), are connected to the Internet 102 through communication links 1310 and to the MPDN server 104 through communication links 1311. Any
10 number of client systems 108, 110 ... 112 may also be connected to the Internet 102, through communications links 118, 120 ... 122, respectively. As also stated above, any of the connected devices may communicate with each other in any of a wide variety of communication techniques (e.g., wireless, electrical, digital, analog, light-based, etc.) and protocols (e.g., static or dynamic IP addresses), and through any number of other devices, as would be understood by one of skill in the art.

In the application contemplated by FIG. 13A, the wireless devices 1306 at times request data, for example, images or text from a web site, that must be converted, translated or otherwise processed by wireless device server systems 1304 before it can be transmitted to, and displayed on, a requesting wireless device. Instead of converting the information, the wireless device server systems 1304 may
15 request that the MPDN server 104 accomplish the data conversion or translation. The device server systems 1304 may then provide to the MPDN server 104 any pertinent information, such as information concerning the requesting device, the nature of the data requested, and the processing needed for the data. The MPDN server 104 may then utilize one or more of the client systems 108, 110 ... 112 to

process the data from block 1302 for transmission to the requesting device. In this way, the wireless device server systems 1304 may off-load burdensome and process-intensive conversion tasks to the distributed processing system of the present invention.

5 It is noted the transmission of processed data to the requesting wireless device 1306 may occur in a variety of ways. For example, the processed data may be transmitted from a client system 108 to the server 104, then to the wireless device server 1304 and finally to the wireless devices 1306. Alternatively, the processed data may be transmitted from a client system to the wireless device server 1304, and then to the wireless devices 1306. Still further, the processed data may be transmitted directly
10 from a client system to the wireless devices.

FIG. 13B provides a basic flow diagram for an embodiment of a data conversion process 1350 according to the present invention. In block 1352, a device, such as wireless devices 1306, requests unconverted, non-translated or non-processed data. In block 1354, a server for the device, such as wireless device server systems 1304, processes the data request and contacts the MPDN server 104. In addition, the content provider or server for the requested data, such as a web site content server, may contact the MPDN server 104. The wireless device server systems 1304 provide all pertinent information to the MPDN server 104, such as the type of calling device, its identification, the relevant data requested, and the conversion to take place. The MPDN server 104 then distributes the data and information concerning the requesting device to one or more client systems, such as client systems 108, 110 ... 112, in block 1356. The one or more client systems then convert, translate or otherwise process the data in block 1358. The converted, translated or processed data is then provided to the requesting device in block 1360. Again, in this way, the device servers may provide a wide range of information without having to provide itself the processing power to accomplish the conversion, translation or
25 processing that is required to transmit or display the data on a requesting device.

As shown in FIG. 13B, the device server or the content server 1304 may communicate data and other pertinent information for a conversion directly to the client systems. For example, the MPDN server 104 may provide access to a group of client systems for data conversion purposes for given
30 periods time (e.g., monthly client group allocations), or may provide identities of groups of client systems that may be used at the time a conversion is needed. Once the identity and allocation of client systems to a particular device server or content server is made, the device server or content server may communicate directly with the client systems. In addition, the device server or content server may

provide directly to a requesting device the identity of the one or more client systems accomplishing the data conversion. As shown in FIG. 13B, the requesting device, therefore, may communicate directly with the client system or systems to provide pertinent information concerning the data conversion requested. The client system may then, for example, directly download the desired content and perform the desired data conversion. It is further noted that in addition to the embodiments described above with respect to FIGS. 13A and 13B, other methods for requesting, processing and providing data to and from the requesting device may be implemented with distributed processing system of the present invention, such as caching processed data for later transmission.

FIGS. 14A and 14B depict example block diagrams of file distribution and data sharing through the network fabric, according to the present invention. In particular, FIG. 14A depicts an Internet data file distribution system 1400 that relies upon client systems to provide local data distribution. FIG. 14B depicts a data file distribution system 1450 that allows for data sharing and rapid transmission of a project or data files through the distributed processing system.

Looking now to FIG. 14A, a block diagram is depicted of a distributed processing system 1400 that provides data transmission caching or other local distribution, according to the present invention. In the embodiment of FIG. 14A, server systems 104 are connected through communication link 114 to the Internet back-bone 1402. The Internet back-bone 1402 represents the very high speed connections that carry data long distances, for example, T3 or fiber optic lines that carry Internet data across the United States. A web site 1404 is connected to the Internet back-bone 1402 through communication link 1406, which represents a geographically local connection. The connection block 1410 represents a geographically remote communications link, such as a POP server, head-end machine, telephone line central office, cell site, etc. This communications block 1410 is connected to the Internet back-bone 1402 with a communications link 1408, which also represents a geographically local connection. A variety of client devices and non-client devices 1412A, 1412B, 1412C, 1412D, 1412E and 1412F may be connected below the connection block 1410. It is noted that interface 1414 represents, for example, a secondary network on which client devices 1412D, 1412E and 1412F are connected, such as a home network.

In the embodiment shown in FIG. 14A, web site 1404 may be desiring to provide content that is in high demand, over a short period of time. An example of such an event is a live Internet multimedia broadcast. For such an event, there may be a huge influx of devices trying to download the content from

the web site 1404 over a short period of time. The web site 1404 may be unable to meet this extremely large demand, requiring the web site 1404 to shut down.

According to the present invention, the web site 1404 may off-load some or all of its data handling requirements by using the distributed processing system of the present invention for data caching. The web site 1404 may contact server systems 104 and request data caching services. The server systems 104 may then identify a local machine, such as client device 1412E, to act as a local distributor of the content for web site 1404. For example, one or more idle client devices that have been identified, as discussed above, may be utilized as local distributor client device 1412E. The local distributor client device 1412E may first download the content and pass it on to other client and non-client devices 1412B, 1412C and 1412D through communication links 1416A, 1416B and 1416C. It is noted that this caching will be aided if the client and non-client devices receiving the cached data are relatively short communication hops from local distributor client device 1412E.

This data or network caching allows data to be streamed to an end user level device, which may then pass the data on to other end user devices. Thus, the downstream communications may be limited, thereby taking the distribution burden off of the web site. For example, web site 1404 may have a large streaming video or multimedia file that is experiencing a heavy load from a given set of network devices. This data file may be cached by a machine, such as client device 1412E, that is below from a communication link 1410. Then, other devices that are also below this communication link 1410 may download the streaming video data from the client device 1412E. This caching eliminates the need to repeatedly send the same data through the same communication links to requesting devices that are located below common communication links. It is noted that the file and data distribution possibilities for this peer file access, caching and data transmission, according to the present invention, are wide and varied and should not be seen as limited to the embodiment shown in FIG. 14A.

FIG. 14B is a block diagram of a distributed processing system 1450 that provides data distribution and data sharing, according to the present invention. As with FIG. 9, FIG. 14B depicts an alternative view of a network fabric that may interconnect any of a wide variety of devices. In the embodiment shown in FIG. 14B, server systems 104 are interconnected with any number of client systems 108A, 108B, 108C, 108D, 108E, 108F, 108G and 108H. Each of the connecting interconnects represents any of a wide variety of communication links that may exist between devices in the network fabric of the present invention. Each of the client systems 108A, 108B, 108C, 108D, 108E, 108F, 108G

and 108H include shared data (SD) according to the present invention. Within this interconnected fabric, block 1452 represents data or project information that is desired to be distributed. The SD blocks within each client system facilitates the distribution of this data or project information.

5 A client agent, as discussed above, installed on the client systems 108A, 108B, 108C, 108D, 108E, 108F, 108G and 108H includes functionality that facilitates a number of services with respect to data transmission and sharing. First, the client agent provides a protected data storage area accessible to outside devices, which is represented by the SD block within each client system in FIG. 14B. This special storage space protects the device from outside devices accessing other storage areas on the device while allowing data to be shared and accessed by other devices and simultaneously used by the local client agent.

10 These shared data (SD) blocks provide mechanisms that enable a wide variety of possible interactions among the client systems 108A, 108B, 108C, 108D, 108E, 108F, 108G and 108H. For example, the data sharing mechanism may provide a space for a cache of other device addresses attached to the network for both communication purposes as well as security purposes. The mechanism may also provide a simple indexing system that is automatically re-indexed when content is added or removed from the storage area. This indexing system may provide a mechanism for other client agents to perform discovery on the local client information and visa versa. Through information stored within this shared data, the distributed processing system of the present invention facilitates many distributed file system applications such as distributed resume posting, distributed caching, distributed advertisement serving, etc. In addition to the above, the storage block (SD) within each client system may include an interface for displaying or playing data types (such as images, audio files, video files, etc.) stored both locally and/or remotely on other client devices. This would enable simple picture sharing, for example, between remote families connected via the internet, as part of being a client system within the distributed processing system of the present invention.

25 In the embodiment shown in FIG. 14B, data or project 1452 is injected into the fabric through a connection to client system 108C and server systems 104. These connections represent that the information may pass first to server systems 104, or may pass first to a client system, such as client system 108C. It is noted that there are other ways that the data may be injected into the fabric. Once injected, the data 1452 may be transmitted throughout the fabric through any of a wide variety of communications, including client-to-client, server-to-client, client-to-server, client-to-non-client, non-

client-to-client communications, and/or non-client-to-non-client communications. These communications may be based upon a variety of mechanisms, such as polling mechanisms and pre-assigned firewall ports. This technique provides a vehicle that facilitates the distribution of information to a large number of devices in a short period of time.

5

Applications for this data distribution are widely varied. For example, any important file that is time sensitive may be propagated to a large number of client devices, non-client devices, servers, or other connected devices, in a short amount of time. This transmission may occur quickly and efficiently once the information is injected into the distributed processing system of the present invention. Example time sensitive data files are anti-virus signature files, which when distributed through the distributed processing system of the present invention, may be transmitted through the network fabric faster than a new virus may normally proliferate.

10

Another application for rapid propagation of files is utilizing this technique for propagation of workloads. One example is distributed resume or job searching. In such a system, participating job seekers and participating employers may rapidly search for one another. A job seeker may inject a job request or search into the fabric that is then routed by each successive device to other devices without the need for control from the server systems 104. Similarly, an employer may inject candidate criteria into the fabric that is then routed to successive devices. The result is an extremely fast search and identification of employers and candidates.

15

20

25

30

FIG. 15 is a block diagram of an alternative representation for a distributed processing system 100, according to the present invention. Server systems 104, database systems 1546 and web interface 1554 are coupled together through communication links 1540, 1542 and 1544. The web interface 1554 includes clients subsystem 1548, task developer subsystem 1550, and advertisers subsystem 1552, and may include other subsystems as desired. The database systems 1546 include workload (WL) information 308, client capability vector information 620, and any other stored information as desired. Server systems include various modules and subsystems, including database interface 1532, web server 1536, task module and work unit manager 1530, client statistics module 1534, advertising manager 1538, task module version/phase control subsystem 1528, sweepstakes engine 1524, server control subsystem 1526, and communication interface 1522. It is noted that in the embodiment of a distributed processing system 100 as depicted in FIG. 15, the three primary operations for the server systems 104, database

systems 1546 and web interface 1554 are directed to managing, processing and providing an interface for client systems, customer tasks, and customer advertising.

As discussed above, each client system includes a client agent that operates on the client system and manages the workloads and processes of the distributed processing system. As shown in FIG. 15, each of the client agents 270A, 270B ... 270C communicates with the server systems 104 through communication links 1516, 1518 ... 1520, respectively. As discussed above, any number of different techniques and architectures may be utilized to provide these communication links. In the embodiment as shown in FIG. 15 with respect to client agent 270A, each client agent includes a base distributed processing system component 1506 and a separate project or workload component 1504. As depicted, a communication interface 1508, a core agent module 1502, and a user interface 1510 make up the base distributed processing system component 1506. The task module 1512 and the work unit 1514 make up the separate project or workload component 1504. The task module 1512 operates on top of the core agent module 1502 to provide processing of each project work unit 1514. It is noted that different or additional modules, subsystems or components may be included within the client agent, as desired. For example, a personal computer screen saver component may be part of the base distributed processing system component 1506 or the separate project or workload component 1504.

Also as discussed above, security subsystems and interfaces may be included to provide for secure interactions between the various devices and systems of the distributed processing system 100. As depicted in FIG. 15, a security subsystem and interface 1560 is interconnected with the server systems 104, the database systems 1546, the web interface 1554, and the client agents 270A, 270B ... 270C. These interconnections are represented by lines 1566, 1564, 1562, and 1568, respectively. The security subsystem and interface 1560 operates to secure the communications and operations of the distributed processing system. This security subsystem and interface 1560 also represents a variety of potential security architectures, techniques and features that may be utilized. This security may provide, for example, authentication of devices when they send and receive transmissions, so that a sending device verifies the authenticity of the receiving device and/or the receiving device verifies the authenticity of the sending device. In addition, this security may provide for encryption of transmissions between the devices and systems of the distributed processing system. The security subsystem and interface 1560 may also be implemented in a variety of ways, including utilizing security subsystems within each device or security measures shared among multiple devices, so that security is provided for all interactions of the devices within the distributed processing system. In this way, for example,

security measures may be set in place to make sure that no unauthorized entry is made into the programming or operations of any portion of the distributed processing system including the client agents 270A, 270B ... 270C.

5 In operation, client systems or end-users may utilize the clients subsystem 1548 within the web interface 1554 to register, set user preferences, check statistics, check sweepstakes entries, or accomplish any other user interface option made available, as desired. Advertising customers may utilize the advertisers subsystem 1552 within the web interface 1554 to register, add or modify banner or other advertisements, set up rules for serving advertisements, check advertising statistics (e.g., click statistics),
10 or accomplish any other advertiser interface option made available, as desired. Customers and their respective task or project developers may utilize the task developer subsystem 1550 to access information within database systems 1546 and modules within the server systems 104, such as the version/phase control subsystem 1528, the task module and work unit manager 1530, and the workload information 308. Customers may also check project results, add new work units, check defect reports, or accomplish any other customer or developer interface option made available, as desired.

Advantageously, the customer or developer may provide the details of the project to be processed, including specific program code and algorithms that will process the data, in addition to any data to be processed. In the embodiment shown in FIG. 15, this program code takes the form of a task module 1512 within the workload, while the data takes the form of work unit 1514. These two portions make up the project or workload component 1504 of each client agent 270. For a given project, the task module 1512 will likely remain relatively constant, except for version updates, patches or phase modifications, while the work unit 1514 will likely change each time processing of the data that it represents is completed. The project or workload component 1504 runs in conjunction with the base
25 distributed processing system component 1506. When a different customer or project is started on a given client system, the project or workload component 1504 will typically be replaced, while the base distributed processing system component 1506 will likely remain relatively constant, except for version updates, patches or other modifications made for the distributed processing system.

30 Information sent from the server systems 104 to the client agents 270A, 270B ... 270C may include task modules, data for work units, and advertising information. Information sent from the client agents 270A, 270B ... 270C to the server systems 104 may include user information, system information and capabilities, current task module version and phase information, and results. The database systems

1546 may hold any relevant information desired, such as workload information (WL) 308 and client capability vectors (CV) 620. Examples of information that may be stored include user information, client system information, client platform information, task modules, phase control information, version information, work units, data, results, advertiser information, advertisement content, advertisement purchase information, advertisement rules, or any other pertinent information.

Now looking to FIGS. 16, 17A, 17B, 18A and 18B, an embodiment for security features for the distributed processing of the present invention will be described. FIG. 16 provides a representation of the distributed processing environment including security subsystems. FIGS. 17A and 17B provide block diagrams of the communication interface between client systems and the server systems. And FIGS. 18A and 18B provide detailed block diagrams of an embodiment of security measures for the server systems and the client systems.

Referring to FIG. 16, an embodiment 1600 of a distributed processing system is depicted. Server systems 104 include a security subsystem 354 through which communications to and from the server systems 104 may be made secure. Client systems 108A, 108B ... 108C and client systems 108D, 108E ... 108F represent any number of client systems that may communicate with server systems 104 or with each other. Each of the client systems 108A, 108B, 108C, 108D, 108E and 108F include a security subsystem 272A, 272B, 272C, 272D, 272E and 272F, respectively. The electronic information 1602 represents information that the server systems 104 is to communicate to client systems 108A, 108B, 108C, 108D, 108E and 108F in a secure manner, so that no unintended or intercepting recipient may understand or tamper with the electronic information 1602, and so that no third party may insert non-authorized information into the distributed processing system 1600. Although not shown, it is understood that any one of the client systems 108A, 108B, 108C, 108D, 108E and 108F may have electronic information that is to be securely sent to the server systems 104 or to any other of the client systems 108A, 108B, 108C, 108D, 108E and 108F.

Electronic information 1602 represents information that is communicated to facilitate the operations of the distributed processing system 1600. Such information includes the client agents that are downloaded to each client system, the workload applications for any given workload, and any work unit that will be processed by a client system. Electronic information 1602 may also be any type of information to be sent or received within the distributed processing system, such as text, images, audio streams, video streams, databases, spreadsheets, PDF files, Shockwave data, Flash data, applications,

data files, chat streams, or any other information, data or data streams. In addition, electronic information may be sent by client systems 108A, 108B, 108C, 108D, 108E and 108F to the server systems 104 and/or any of the other client systems.

5 The Certificate Authority (CA) block 1604 within the server systems 104 represents an entity that helps to ensure validity of encryption and decryption codes. For example, within a public/private key encryption environment, a Certificate Authority may help ensure that a public key alleged to be from a particular entity is in fact legitimately from that entity. One third-party entity that performs this CA function on the Internet is Verisign, Inc. Having a third-party perform the CA function can be
10 advantageous in a transaction or communication between non-trusted entities. For example, the sending entity provides its public key information to the third-party CA, which verifies the information and creates a certificate that includes the sending entity's public key information. This certificate may then be encrypted and signed by the third-party CA. The receiving entity may then obtain the certificate from the third-party CA and decrypt it with the third-party CA's public key. The receiving party will then have the sending party's public key and be fairly secure that it is a legitimate public key from the sending party.

As shown in FIG. 16, the CA functionality may be part of the server systems 104, such that the server systems 104 act as their own Certificate Authority with respect to client systems 108A, 108B, 108C, 108D, 108E and 108F and any other devices that are part of the distributed processing system. A third-party CA is not as needed in this distributed processing environment because the server systems 104 primarily direct the operations of the distributed processing system. Thus, there is less of a need for a third-party entity to provide a CA function. It is noted that CA functionality may be provided only by the server systems 104, only by third-party CAs, or any combination of server systems 104 and third
25 party CAs, as desired for a particular embodiment. In addition, if desired, no CA functionality could be provided so that secure communications between the server systems 104 and the devices within the distributed processing system were conducted without the use of a Certificate Authority.

FIG. 17A is a block diagram of an embodiment 1700 for a communication interface between a
30 client system 108 and the server systems 104. In this embodiment 1700, the network is preferably the Internet. As depicted, the client system 108 includes a client agent 270 and a network browser 1702. The server systems 104 include a client agent download site 1710, from which the client system 108 may download the client agent 270 through communications 1704. The server systems 104 also include block

1712, which represents a variety of client service functions that may be provided by the web interface for the server systems 104 through communications 1706. For example, in a public/private key security environment, a client system 108 may download from block 1712 a Certificate Authority (CA) certificate that includes the server public key. In addition, the client system 108 may login to the web page interface for the server systems 104. And the server systems 104 may generate dynamic certificates. The client system 108 may also send and receive information to application server 1714 through communications 1708, for example, to receive project work units. Finally, as depicted, database systems 1546 may send information to and receive information from the blocks 1710, 1712 and 1714 of the server systems 104 through communications 1716, 1718 and 1720. As discussed more above, database systems 1546 may include any desired information, for example, a workload database 308 and/or a capability vector database 620.

FIG. 17B is a block diagram for an Internet communication protocol structure 1750 that may be utilized for communications 1704, 1706 and 1708. As depicted in FIG. 17B, three basic application layers are utilized by each client system 108 and the server systems 104 to communicate with each other. The TCP/IP layer 1756 represents a standard Internet communication protocol that allows devices to identify and send information to each other across the Internet, as is well known to those of skill in the art. The secure network layer (SNL) 1754, such as the secure socket layer (SSL), represents a protocol that allows devices to confirm the identity of servers and the other devices with whom they communicate, as long as those servers or other devices utilize similar protocols. The application security level 1752 represents other desired security or communication protocols that are implemented by programs running on the client system 108 and/or the server systems 104.

In operation, the server systems 104 may secure the download of the client agent 270 to the client system 108 by requiring that the client system 108 download the client agent 270 from the client agent download site 1710. As part of the server authentication sequence, the download site 1710 will send back an identifier to assure users that they are indeed connected to the proper server systems 104. This identifier may be, for example, a CA certificate, but may be any other identifier, as desired. Because it is desirable to have the client agent running on as many distributed devices as possible for the distributed processing system of the present invention, user authentication may not be required to download the client agent 270 from the download site 1710.

09/24/05 05:47:04
FOI b7E b7D

Once a client system 108 has downloaded and installed the client agent 270, the client system 108 will communicate with the application server 1714 to begin working within the distributed processing system. For these communications, server and client authentication may be required to help ensure security. To accomplish this authentication, for example, two-way authentication may be utilized.

5 To provide a public/private key combination for the client agent 270, each client agent 270 that is downloaded by a client system 108 may have embedded within its code a default identifier and a default public/private key pair. Thus, the server systems 104 may use secure network protocols (such as SSL or similar schemes) to authenticate each client system 108, and each client system 108 may use compatible protocols to authenticate each server application with which it communicates. These applications, for

10 example, may include the functionality provided by blocks 1712 and 1714, and, therefore, the communications 1706 and 1708 would utilize authentication.

As an alternative to embedding a public/private key combination and associated identifiers or certificates into the client agent 270, the public/private key pairs may be dynamically generated in block 1712. For example, at start-up, at reboot or at some desired time or event, the client system 108 may generate a new public/private key pair. When the client system 108 next communicates with the server systems 104, the client system 108 request a certificate from the server systems 104. The server systems 104 may then act as a Certificate Authority (CA) and provide a CA certificate to the client system 108. This dynamic certificate generation, therefore, allows for added security by allowing each client system 108 to have its own public/private key pair for secure network protocol communications and by having this key pair change at some desired recurring event for the client system 108, such as reboot.

The client system 108 may initiate its communication with the server systems 104 by logging on to the authentication server, which may be part of block 1712. The user may be prompted to enter a valid e-mail address and/or password, if already registered, or may be asked to register if the e-mail address and/or password are not recognized. Once registration is completed, a password may be e-mailed back to the user to provide validation of the user. If authentication is successful when a user logs into the server systems 104, the server systems 104 may provide a host-ID, and user-ID and a session key for any given communication session.

It is also desirable that once a user has successfully registered, the user may install the client agent 270 on any number of other host or user systems without needing to interact with that systems network browser, other than to set host-specific preferences. For example, when downloaded, the client

agent 270 may take the form of a self-extracting program that installs the appropriate files to the client system 108, including the proper host and user identifications. In addition, to help ensure proper identification, the session keys may be exchanged each time the client system 108 communicates with the server systems 104. For example, the client system 108 may communicate its current session key to the server systems 104 each time it communicates with the server systems 104. The server systems 104 will then send a new session key for the client system 108 to utilize for the next session. In this way, stale identification information may be reduced. In addition to this security feature, communications may also be encrypted and decrypted with various encryption techniques, as desired.

Referring now to FIGS. 18A and 18B, one embodiment will be discussed for a security model utilizing public/private key encryption. This security model utilizes a third-party CA to provide a CA certificate for the server systems 104.

FIG. 18A is a block diagram of an embodiment 1800 for security procedures implemented by server systems 104. Electronic information 1602 is to be communicated to a client system 108. This electronic information 1602 travels through four different paths that provide security information.

One path begins with the electronic information 1602 being encrypted with the server private key in block 1802. Then, in block 1830, the encrypted information is sent to client systems. This encrypted information is represented by arrow 1826.

A second path flows from block 1802 to block 1804 where a hash value is generated for the encrypted electronic information. It is noted that a hash value is a unique value that may be generated for any given electronic file based upon the contents of that file and the algorithm used to calculate the unique value. There are any number of algorithms that may be used to calculate a hash value, as would be understood by one of skill in the art. Proceeding down the second path to block 1806, the hash value generated on the server side for the encrypted electronic information (i.e., the information sent to the client system from block 1830 via 1826) is compared with a hash value 1822 from the client system 108. This hash value 1822 represents the client system's calculation of the hash value for the encrypted electronic information that the client system 108 received from the server system 104. If no tampering has occurred and the data was transmitted accurately, the client system hash value should match the server hash value. In block 1808, the server systems 104 provide an indication of the result of the hash check evaluation back to the client system 108. This pass/fail determination is indicated by arrow 1824.

5 A third path begins with block 1810 where a hash value is calculated for non-encrypted electronic information 1602. This hash value is then encrypted in block 1816 with the server private key. Next, this encrypted hash value is sent to the client system 108 in block 1818. The arrow 1821 represents the encrypted hash value for the non-encrypted electronic information.

10 A fourth path, and the last depicted in the embodiment 1800 of FIG. 18A, flows from block 1810 to block 1812, where the hash value is partitioned into N different portions. These N different portions are preferably designated for N different client systems 108, as well as any client systems 108 receiving a redundant distribution of any one of the N different portions. In block 1814, the N different hash value portions are encrypted with the server private key. Next, the N different encrypted hash value portions are sent in block 1820 to N different client systems 108, as well as being sent to client systems 108 receiving redundant distributions of the hash value portions. The arrows 1828 represent the distribution of the N different hash value portions. It is noted that redundant distribution of the N hash value portions is desirable because, as discussed below with respect to FIG. 18B, when the hash value is reconstructed by a client system 108, it is desirable to have multiple sources for each portion in case one of the receiving client systems is not available at any given time.

25 Looking now to FIG. 18B, the corresponding security procedures implemented by a client system 108 are discussed with respect to embodiment 1850. Initially, in block 1854, the client system 108 receives CA certificate 1852 containing the server public key and the server identity. It is again noted that other unique identifiers may be utilized instead of CA certificates, as described above. If a CA certificate is utilized, this CA certificate may be provided from a third-party Certificate Authority (CA) or from the server systems 104 or any other desired source. In block 1856, the client system 108 verifies the accuracy of the CA certificate using the CA's public key. If this verification is not successful, the client system 108 may wait some period of time before retrying. In addition, the time period may be a random period of time. In addition, as discussed with respect to FIGS. 17A and 17B, the client system 108 will login to the server systems 104. If this authentication is not successful in this login, the client system will notify the user of the system and the server systems 104, and then wait for some period of time or a random amount of time before attempting to re-verify.

30 In block 1862, the client system 108 receives the encrypted information 1826. Next, the client system 108 creates a hash value for the encrypted information in block 1864. This hash value is

0934785-01301
FOI b7D b7E b7F

preferably calculated using the same algorithm utilized by the server systems 104 in generating the hash value for the encrypted information in block 1804 of FIG. 18A. Once the client system 108 has calculated the hash value for the encrypted information, this hash value 1822 is sent to the server systems from block 1866. As discussed above, a pass/fail response 1824 is sent back by the server systems 104.

5 This hash check evaluation is received in block 1868. If the check was a FAIL, flow passes to block 1870 where the client system 108 sends out a notice to the server systems 104 and any other client system to which it is attached that a problem has been encountered. The client system 108 then ends the current connection with the server systems 104. It is noted that the client system 108 may retry several times before moving onto block 1870, and that the reporting scheme may be modified, altered or

10 developed as desired.

If the hash check evaluation was a PASS, flow passes to block 1872 where the electronic information is decrypted with the server public key, which was verified in block 1856. A hash value is then calculated for the electronic information 1874. Again, the hash generation algorithm is preferably the same as that used by the server systems 104 in creating the hash value in block 1810 of FIG. 18A. Next, the hash value is sent from block 1874 to block 1886, where it is compared with two other hash value calculations.

One of the other hash values comes from a path that begins with block 1858, in which the client system 108 receives the encrypted hash value 1821 for the non-encrypted information. In block 1860, the encrypted hash value is decrypted with the server public key. The hash value is then sent to block 1886.

The third hash value for block 1886 comes from a path that utilizes the N different hash portions sent out by the server systems in block 1820 of FIG. 18A. In block 1876, the client system receives a portion 1828A of the partitioned hash value 1828. In addition to one of the partitioned hash values, it is noted that the server systems 104 will also send information providing the identity and source for the N-1 other hash value portions. In block 1878, the client system 108 decrypts the portion 1828A with the server public key. Next, in block 1880, the client system 108 resolves the identity of the source for the

25 N-1 other portions, which may be N-1 other client systems. In block 1882, the client system 108 obtains the N-1 other portions, and assembles the N partitions into a hash value for the non-encrypted electronic information in block 1884. The resulting hash value is then sent to block 1886. It is noted, as indicated above, that redundant distribution of the N portions of the partitioned hash value is desirable so that

30

unavailability of one client system will not cause another client system to be unable to re-assemble the N different portions.

Once the three hash values are received in block 1886 from three different sources, they are compared to see if they match. If this check is a FAIL, flow moves to block 1888, where the client system 108 sends out a notice to the server systems 104 and any other client system to which it is attached that a problem has been encountered. The client system 108 may also inform the client systems from which it received the N-1 other portions, and the client system 108 may retry the procedures, if desired. In addition, once a client system 108 is notified of a potential problem, the client system 108 may download a special check file from the server systems 104 to make sure that the server systems have not been compromised. If still a FAIL, the client system 108 then ends the current connection with the server systems 104. If the check is a PASS, the electronic information is utilized, as represented by block 1890.

FIGS. 19 and 20 provide block diagrams for further describing the distributed processing system and environment of the present invention that allows for third parties, such as network service providers, to monetize, or gain revenue from, their respective user bases.

Looking first to FIG. 19, a block diagram is depicted for a distributed processing system 100 and environment 1900 in which network service providers are enabled to monetize their user bases. Environment 1900 includes a distributed processing system 100, a customer 152, and a third-party network service provider 1902. The customer 152 represents an entity that has a project 1912 that the customer 152 would like processed by the distributed computing system 100. In return for the processing of the project data and the results 1914 of this processing, the customer 152 will often make a payment 1916. The third-party network service provider 1902 maintains a user database 1904 that identifies its user base 1920 including users 1906A, 1906B ... 1906C.

The service provider 1902 may be, for example, an Internet business that provides any of a variety of services to users, such as Internet access, e-mail, web page hosting, domain name hosting, file sharing services or any other Internet-based service. In addition, such Internet-based services may be offered for free or low cost to users, in which case the users have historically agreed to view banner or other advertisements in return for the free or low cost service. However, as stated above, advertising revenue has been subject to diminished pricing and has become a unreliable source of revenue for many

Internet-based companies. To facilitate the number of projects that the distributed processing system 100 can take on and the speed at which these projects can be processed and completed, it is desirable to increase the amount and capabilities of the computing resources available to the distributed processing system 100. To the extent that the users of the service provider 1902 represent a pool of underutilized resources, these users represent a potentially valuable resource to the distributed processing system 100.

According to the present invention, the service provider 1902 may realize value from its user base and thereby monetize this user base by facilitating the use by the distributed processing system 100 of computing resources related to these users. Thus, for example, in return for free services, the users may agree to have their respective computing resources utilized by the distributed processing system 100. The service provider 1902 may then provide to the distributed processing system 100 the user identifications (IDs) 1908 related to its user base in return for revenue sharing 1910. This monetizing architecture according to the present invention thereby provides a significantly advantageous avenue for service providers or other entities that control or have user bases with useful processing capabilities, such as Internet-based service providers, to generate revenue from its user base, particularly in the face of falling revenue from other sources, such as advertising revenue.

The revenue sharing 1910 may be, for example, a share of payment 1916 relative to the amount of processing toward the project 1912 that was completed through the use of the user resources 1922 made available through users 1906A, 1906B ... 1906C. It is noted that the revenue sharing 1910 may take any desired form, including but not limited to (a) upfront payments based upon attributes of the user base, such as size or processing capabilities, (b) payments based upon the number of users that become members of the distributed processing system, (c) payments based upon the types of projects processed by the user base, or (d) any other desired compensation scheme related to the value of the user base being made available by the third party.

The monetizing invention, therefore, focuses on capabilities of internet, intranet, wireless or otherwise network connected PCs, internet appliances, notebook computers, servers, storage devices, NAS (Network Attached Storage), or any other connected computing device that could provide any of a number of useful capabilities and that is part of a underutilized user base, such as user bases of Internet-based businesses that rely on advertising or any other method of monetizing their user base in exchange for a valuable service (e.g. free internet access, email, etc.). As discussed above, these useful processing capabilities span the entire range of generic computing subsystems including: Central Processing Unit(s)

(CPUs), Digital Signal Processor(s) (DSPs), Graphics Processing Engine(s) (GPEs), Hard Drive(s) (HDs), Memory (MEM), Audio Subsystem(s) (ASs), Communications Subsystem(s) (CSs), Removable Media Types (RMs), or other Add-In/On Accessories (A/OAs) with potentially useful unused capabilities. Market creation and potential compensation for all unused capabilities can be accomplished through the massively parallel distributed software architecture of the distributed processing system 100. For example, credits (revenues) would be generated each time a unit of work is accomplished by one (or more) of the subsystems on a user's computing device via a client agent installed on the device for the purposes of managing, processing and completing units of work. The total credits/revenues generated may be, for example, dynamic depending on how many are received. Through this architecture of the present invention, significant revenues may be generated from the user base of the service provider where the service provider may have previously been unable to monetize effectively this user base.

It is further noted in more general respects that the entity 1902 may be any entity that has direct or indirect control over a group of users, such that the users resources may be offered to and utilized by the distributed processing system 100. An example of one such more general entity would be a company that has a large group of internal users that are capable of being linked to the distributed processing system 100, for example, through an intranet network of computer systems or computing devices. The computing resources related to these users may also be monetized according to the present invention.

Looking now to FIG. 20, a block diagram for an embodiment 2000 representing the components for a client agent 270 along with a representative indication of responsibility for those components. Client agent 270 includes a core agent component 1502, a project component 1504 and a user interface component 1510. As discussed above, the core agent component 1502 can provide the base distributed processing functionality for the client agent 270. The project component 1504 can provide the project-specific functionality for the client agent 270. And the user interface 1510 can provide any desired user viewable information or interaction functionality. These three general components may be modular software components such that the project component 1504, the core agent component 1502, and the user interface component 1510 may be separate software modules, if desired, that link together through appropriate APIs (Application Programming Interface). Thus, each of these components may be designed and developed independently or jointly, as desired. In effect, the core agent component 1502 can provide a backbone upon which is attached the project component 1504, the user interface component 1510, and any other desired component. Thus, when a new project or interface is desired for any given client agent 270, for example, this component may be efficiently replaced with the new

component in a modular fashion still utilizing the core agent component 1502 as the backbone. In addition, each component may be updated and modified without requiring modification or updates to the other component software code.

Also depicted in FIG. 20 are customer 152, distributed processing system (DPS) 100 and service provider 1902, which are in communication with each other through interactions or interfaces 2012 and 2014. In this embodiment of FIG. 20, the customer 152 is represented as providing the software development and code 2002 for the project component 1504. The distributed processing system 100 is represented as providing the core agent code 2008 for the core agent component 1502. And the service provider is represented as providing at least a portion of the interface development and code 2010 for the user interface component 1510. In operation, the workloads 2004 and the results 2006 would still typically be under the control of the distributed processing system 100.

It is noted that this modular architecture facilitates the development of project software code and interface software code by entities other than the owner of the distributed processing system 100. For example, with respect to FIG. 19, an Internet-based service provider may have a user interface already designed and implemented for its user base, such as a web browser user interface for users of free Internet access services provided by such a service provider. Once the core agent component 1502 is installed on a user's computer, the existing third-party user interface may hook into the core agent component 1502, thereby making the user's resources available to the distributed processing system 100, while maintaining the user interface the user has come to expect from the service provider. Thus, the service provider 1902 may provide the user interface it desires for the service it is providing, while at the same time monetizing its user base by facilitating its users becoming part of the available resources for the distributed processing system 100.

FIG. 21 provides a block diagram for an embodiment 2100 for a distributed processing system that provides network attached storage (NAS) services, according to the present invention. Although the distributed computing, client agent, file sharing and distributed data storage functionality, as described above, provides data storage that is connected to a network, the NAS embodiment 2100 depicted in FIG. 21 is provided herein for additional explanation. In particular, standard NAS device functionality provided by typical dedicated NAS devices can be accomplished according to the present invention by utilizing unused or under-utilized resources on network-connected devices through capabilities provided by the client agent running on those devices. The resulting software-based NAS services hosted on a

distributed computing platform, according to the present invention, thereby provide an advantage storage solution.

As with other embodiments described above, the NAS embodiment 2100 focuses on taking advantage of unused or under-utilized capabilities of Internet, intranet, wireless or otherwise network connected personal computers (PCs), internet appliances, notebook computers, servers, storage devices or any other connected computing device. Such devices have typically been put into operation for purposes entirely separate from operating as a part of a distributed processing or storage system and have a primary operational purpose that is not to provide NAS services. Through the client agent and infrastructure implementation of the present invention, such devices can provide storage capabilities that allow these devices to appear to users as dedicated NAS devices, while still being used for the purpose for which they were put into operation. For example, a large number of desktop PCs in the enterprise (Intranet) or on the Internet can each act as a NAS device by running a client agent program and NAS component that brings its resources to the network with the appearance of a dedicated NAS device or as part of an integrated system that appears as a single or dedicated NAS device. The implementation can be accomplished in a variety of ways. A relatively simple implementation is to have the devices simply mimic the functionality of stand-alone NAS devices, while a more complex implementation is to provide a storage fabric across the Intranet or Internet where hundreds or thousands of devices work cooperatively as NAS devices in conjunction with server systems that assist in controlling and managing the storage function of the NAS fabric.

Looking now to FIG. 21, the NAS embodiment 2100 includes a network 102, which as discussed above, can be any network communication fabric that provides connectivity between various machines. As shown, connected to this network 102 are a user device fabric 2102, a client device fabric 2104, and servers systems 104 and 154D. More particularly, the user device fabric 2102 may include any number of potential user devices 152A, 152B ... 152C. The client device fabric 2104 may include any number of potential client devices 108A, 108B ... 108C. It is noted that that the user devices 152A, 152B ... 152C may be configured to communicate directly to each other, the client devices 108A, 108B ... 108C may be configured to communicate directly to each other, and the user devices 152A, 152B ... 152C and the client devices 108A, 108B ... 108C may be configured to communicate directly between themselves without intervention by the server systems, if desired. The servers systems may include user server systems 154D, which have attached database systems 2112, and NAS server systems 104, which have attached NAS database systems 2108 and a centralized data location library 2110. To provide NAS

functionality, the client devices 108A, 108B ... 108C run a client agent, as described above, that provides NAS functionality, for example, by running a NAS component as one of the projects for the distributed device. In addition, some of these NAS devices, such as NAS client devices 108A and 108B, may store distributed data location information, such as data location libraries 2106A and 2106B.

5

The NAS distributing computing embodiment 2100 depicted provides an environment in which a variety of different NAS configurations may be utilized, as desired. These configurations include, for example, at least three different NAS architectures that may be utilized: a stand-alone only NAS device implementation, a NAS device fabric implementation, and a server assisted implementation. It is noted that variations, combinations and hybrids of the these implementation examples may be utilized, as desired, without departing from the present invention.

10

First, the NAS client devices within the client device fabric 2104 can act as stand-alone NAS devices, and the user devices within the user device fabric 2102 can directly access them as such. In this first implementation, the NAS devices 108A, 108B ... 108C run client agent programs that configure the devices to be capable of mimicking dedicated stand-alone NAS devices. Each of these software-based NAS devices is self-contained, knows nothing of other software-based NAS or other NAS devices, and operates like a standard dedicated NAS device, except for the fact that the software-based NAS devices 108A, 108B ... 108C are actually non-dedicated devices that rely on spare and/or unused capabilities of the existing devices to provide NAS services. In this stand-alone only implementation, the data location libraries 2106A and 2106B would not be required to be utilized.

09/24/2015 10:43:01
FOE-F-054736

Second, at least some of the NAS client devices within the client device fabric 2104 include data location information or libraries that are used to provide links or directions to data distributed on the various NAS client devices 108A, 108B ... 108C. These links help provide a NAS device fabric from which user server systems 154D and user devices 152A, 152B ... 152C within the user device fabric 2102 can draw storage resources. For example, in the embodiment shown in FIG. 21, NAS client devices 108A and 108B can have data location libraries 2106A and 2106B, respectively. These libraries provide information to a user about where desired data is stored across the entire NAS client device fabric 2104, for example, by providing a distributed database index of the files and other data being stored on the myriad of software-based NAS devices. This database index, or portions thereof, can be stored within each data location library and can be organized as desired to provide direction information for users. In this NAS device fabric implementation, the NAS client devices 108A, 108B ... 108C

25

30

provide both storage and data location functionality, thereby acting both to provide data storage services and distributed data location services. The NAS devices storing the data location information, for example, would receive storage and access requests from user devices and would direct the user devices to the distributed devices storing the requested data. This NAS device fabric implementation would again
5 rely upon client agent programs running on the client devices and would not require the use of centralized server resources to enable the system to perform the NAS services. It is noted that not all software-based NAS devices need to act as both stores of information and location information hosts, such as NAS client device 108C, as depicted in FIG. 21.

10 In the third server assisted architecture, server systems 104 facilitate the NAS functionality, for example, by providing centralized management, control and support for the NAS architecture. To this end, the NAS servers systems 104 may include, for example, file storage database systems 2108 and a centralized data location library 2110. In this implementation, therefore, the NAS client devices 108A, 108B ... 108C that are running client agents which make them capable of mimicking dedicated NAS devices are combined with centralized database and server systems 104, 2108 and 2110 that provide services, such as indexing, monitoring, back-up, security and other services for the files and other data being stored on the multitude of software-based NAS devices creating the NAS fabric. Thus, the NAS server systems storing data location information in the centralized data location library 2110, for example, would receive storage and access requests from user devices and would at least in part direct the user devices to the distributed devices storing the requested data. It is noted that the server systems 104 could direct the users to the requested data using only the centralized data location library 2110 or with assistance from other resources such as the data location information stored by the NAS devices.

25 It is further noted, as also discussed above, that the NAS server systems 104, as with the user server systems 154D, are typically a single server machine or a relatively small group of server machines that are commonly controlled and managed, for example, as to physical location and authorized use. In contrast, the user device fabric 2102 and the client device fabric 2104 are typically a multitude of devices that do not have clear bounds. For example, with respect to Internet-connected user devices and client devices, these devices are typically not related and not commonly controlled or managed. With respect
30 to intranet-connected user devices and client devices, there will often be more common control and management, for example, with respect to internal company computer resources. As contemplated by the present invention, once the software-based NAS functionality is implemented, any of a variety of

09634785-01301
FILED

user devices 152A, 152B ... 152C and/or user server systems 154D may take advantage of the software-based NAS resources made available by the distributed computing system of the present invention.

5 In providing NAS services, it is advantageous for the software-based NAS implementation of the present invention to remain unobtrusive to the end-user or owner of the device from a storage perspective. It is also advantageous to facilitate the full use of available storage space. To help achieve and ensure these objectives, the NAS implementation may undertake a number of controls to minimize interference with user storage operations while maximizing the storage value of the software-based NAS solution. These controls may include user or server controlled management features and may include, 10 for example, user specified constraints, standard bit, block or file priority levels and direct bit or file priority markings. Using such storage priority controls, the present invention may maximize storage space efficiency by helping to take advantage of the storage space so that it is operated at a full level. It is noted that these storage priority controls may also be utilized for activities and operations, other than NAS services, that are being provided by a client agent program running on a distributed device.

With respect to the user specified constraints, the client agent program running on a device would recognize and adhere to user-based constraints that may be specified by the user. These constraints may be set by a user, owner, administrator, etc. for the device and may include settable parameters such as a maximum allowed percentage for use of existing storage space, a maximum number of bytes that may be used (MB, GB, TB, etc.), or a minimum amount of storage space that must be left free within the available storage. When these user-based constraints are utilized, the NAS component will recognize the constraints placed on it and adhere to the constraints, up to and including, deleting all shared storage on the device and becoming dormant.

25 With respect to storage priority levels, the client agent program may rely upon priority levels for bits on the storage device, blocks of data on the device, files on the device, or fragments of files on the device to determine and facilitate how data storage is managed. By using storage priority levels, any of a wide variety of schemes may be implemented for handling storage and storage conflicts and these schemes may be selected depending upon the circumstances for any given client system. In this way, 30 when space on the device is running low, older or less important data or files can be deleted according to fixed or dynamic algorithms. For example, file or data assessment and storage management schemes could include, but are not limited to, a deletion scheme that uses an algorithm, such as one of the following:

09034735-04301

- Timing of storage, for example, straight first-in-first-out (FIFO), last-in-last-out (LIFO) or other time/order based algorithm.
- Usage of the files, for example, least-recently-used (LRU), most-recently-used (MRU), etc. cache schemes.
- File content, for example, content based analyses that base decisions on actual file content, such as whether the file or its data is replicated elsewhere in the system, whether the file or its data can be retrieved from other sources outside the current NAS implementation (tape back-up, Internet, etc.), whether the file is a music, video or audio file that can be retrieved from the Internet, etc.
- Duplicate data, for example, back-up set order, age, etc. for backed up files where needed space could be obtained by deleting older backup files or sets where more recent backup files or sets exist.
- Loss-robustness, for example, losing a few bits of data from an audio stream or video stream would not materially damage the retrieval of the data type.

With respect to direct bit, block or file priority marking, the files or data could include a direct marking for QoS (quality of service) levels that would allow efficient management of storage. By using such a direct marker, a file system priority level may be achieved that resembles execution thread priorities in modern operating systems that help fully utilize the processing power of a central processing unit (CPU). Thus, as each program running on a computer may have a CPU-related service level designation, each file or set of data stored on a storage device may include a storage priority level designation. The level of service provided to any given file or set of data by a control and management scheme may then be determined by the storage priority level designation for that file or set of data.

Further modifications and alternative embodiments of this invention will be apparent to those skilled in the art in view of this description. It will be recognized, therefore, that the present invention is not limited by these example arrangements. Accordingly, this description is to be construed as illustrative only and is for the purpose of teaching those skilled in the art the manner of carrying out the invention. It is to be understood that the forms of the invention herein shown and described are to be taken as the presently preferred embodiments. Various changes may be made in the shape, size and arrangement of parts. For example, equivalent elements may be substituted for those illustrated and described herein, and certain features of the invention may be utilized independently of the use of other

features, all as would be apparent to one skilled in the art after having the benefit of this description of the invention.

09834785-044301